A COMPARISON OF BEARING MEASUREMENTS TO SURFACE FEATURES GENERATED USING STEREOPHOTOCLINOMETRY AND SURFACE FEATURE NAVIGATION TECHNIQUES

Andrew Liounis[*], Coralie Adam[†], Peter G. Antreasian[†], Benjamin Ashman[*], Olivier Barnouin[‡] Jeroen Geeraert[†], Kenneth Getzandanner[*], Chris Gnam[*], Dolan Highsmith[§], Dante S. Lauretta[¶], Jason M. Leonard[†], Joshua Lyzhoft[*], Leilah McCarthy[†], Michael C. Moreau[‖], Derek Nelson[†], and Jeffrey Small[*]

The Origins Spectral Interpretation Resource Identification Security Regolith Explorer (OSIRIS-REx) mission to the asteroid Bennu completed successful two-and-a-half year proximity operations in May 2021. The mission comprehensively mapped Bennu at unprecedented detail and collected a sample of Bennu's surface to return to Earth. Throughout proximity operations, the OSIRIS-REx navigation team used the maps made of Bennu's surface to navigate in the Bennu environment with high accuracy through the use of precise and accurate optical navigation data, radiometric data, and force modelling. The primary type of optical navigation measurements extracted from the images captured by OSIRIS-REx (particularly after first entering orbit around Bennu) were observations of known features on Bennu's surface. Two related but different techniques/tools were used to extract these observations from the images: the Goddard Image Analysis and Navigation Tool Surface Feature Navigation (GIANT SFN) and Stereophotoclinometry (SPC) Autoregister. In this paper we compare the differences between the observables extracted using GIANT SFN and SPC Autoregister, explain the differences, and discuss where each technique is best suited.

INTRODUCTION

The OSIRIS-REx navigation team consisted of a prime team, responsible for creation and delivery of all navigation products to the project, and an independent verification and validation (IV&V) team, responsible for providing support to the prime team as well as independent navigation solutions used by the prime team for consistency checks. To maximize the independent nature of the IV&V team, when performing analysis they used unique software and techniques whenever possible. This was particularly the case in orbit determination (OD) solutions, where the IV&V team used both independent OD software and independent software for extracting measurements to feed the OD software. One of the primary measurement types used during OD for OSIRIS-REx at Bennu involved extracting bearing measurements to known surface features on the surface of Bennu observed in a camera.[1] These measurements are generated by identifying the location of the surface

[*]Navigation and Mission Design Branch, NASA/GSFC, 8800 Greebelt Rd, Greenbelt, MD 20771, USA.

[†]Space Navigation and Flight Dynamics Practice, KinetX Inc., 21 W. Easy St., Ste 108, Simi Valley, CA 93065, USA.

[‡]The Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Rd, Laurel, MD, 20723, USA.

[§]Aerospace Corporation, Chantilly, VA, USA.

[¶]Lunar and Planetary Laboratory, University of Arizona, 1415 N 6th Ave, Tucson, AZ 85705, USA.

[‖]Space Science Mission Operations, NASA/GSFC, 8800 Greebelt Rd, Greenbelt, MD 20771, USA.

features (here defined as a small patch of surface that may contain many typical "features") in the image by predicting what the feature should look like based on the knowledge of the feature topography, a photometric model, and the *a priori* knowledge of the relative position and orientation of the camera with respect to the feature.

To extract these measurements from the image, the prime navigation team used SPC Autoregister.[2] This technique works by (1) predicting the illumination of each feature, (2) extracting the intensity values from the image currently being considered, (3) correlating the predicted and extracted feature illuminations using 2D normalized cross correlation, (4) transforming the shift that best aligns the predicted and extracted feature illuminations from the feature surface to the image frame in units of pixels, (5) using observations of multiple features in the same image to solve a Perspective-n-Point (PnP) problem to correct errors in the *a priori* position and orientation of the camera, and, finally, (6) repeating steps 1–5 until convergence. Importantly, we note that throughout, SPC is operating not on the flat image plane, but rather on the actual surface topography of the feature being considered. Therefore, for each facet of the feature, SPC predicts the intensity value, extracts the observed intensity value for that facet from the image, and then performs the 2D cross correlation in feature space. This is illustrated further in Figure 1.

To remain as independent as possible, the IV&V team developed an alternative method to extract these observables from the images. This alternative technique, termed surface feature navigation (SFN) and implemented in the Goddard Image Analysis and Navigation Tool (GIANT),[3–5] works by (1) rendering what the surface feature is predicted to look like in the image using a single bounce ray trace, (2) correlating between the predicted image and the actual image using 2D normalized cross correlation, (3) solving a PnP problem using observations of multiple features in the same image to correct errors in the *a priori* relative orientation and position, and (4) repeating steps 1–2 with the new estimated relative position and orientation. Everything is performed in the image frame itself, instead of on the surface. In addition, shadows are included in the correlation surface computation, which is generally not the case in SPC.

In this paper we provide a theoretical and numerical comparison of the techniques, discussing both the differences in terms of process, mathematics, and assumptions, as well as differences in the results. When comparing the results of the two techniques, we provide each tool with the same surface features and the same images and compare the observed values. We also compare the consistency of the observed values based on the PnP solutions created by each tool, as well as which features are successfully identified by each technique and which are not. By performing this analysis, we discover the relative strengths and weaknesses of each technique.

THEORETICAL COMPARISON OF THE METHODS

In this section we present a brief theoretical comparison of the SPC Autoregister and GIANT SFN techniques, focusing on where the techniques differ, largely following the discussion in Reference 4. For a more in depth discussion, we refer the reader to Reference 4 and the works cited therein.

Background of SPC Software

SPC Autoregister creates bearing measurements to known surface features through normalized cross correlation[6] of brightness information extracted from an image and projected onto the feature topography against a rendered template representing the same topography. The cross correlation routine is standard, except that it ignores template and image pixels labeled as potential shadows.

Therefore, we focus our discussion on the steps SPC Autoregister takes to render the predicted template and to extract and project the image data onto the template topography.

Generating the Predicted Image To generate the predicted template of the surface, SPC Autoregister computes the expected intensity of each pixel of the feature surface. This is done using a bidirectional reflection distribution function (BRDF) along with the *a priori* viewing geometry at the time the image was captured. SPC Autoregister makes multiple BRDFs available to the user; however, the one most commonly selected, and used in this analysis, is based on the McEwen lunar model, which is a phase-dependent linear combination of the Lambertian and Lommel-Seeliger models:[7]

$$I_k = a_k \left((1 - \beta) \cos(i_k) + \beta \frac{\cos(i_k)}{\cos(i_k) + \cos(r_k)} \right) \tag{1}$$

where for every point k, I is the intensity, a is the albedo, and $\beta = \exp(-\alpha/\alpha_0)$ is a weighting term between the Lambertian and Lommel-Seeliger reflectance models (where α is the phase angle in degrees, and $\alpha_0 = 60°$). Figure 2 shows the viewing geometry used to calculate these angles. We note here that the calculated intensities for each point on the feature remain locked to the feature topography where later correlation with the extracted image will take place.

Generating the Extracted Image To generate the extracted image, SPC Autoregister projects the image intensity values onto the feature topography and then applies a filter to ignore topography pixels that should be invisible to the camera. When projecting the image intensity values, the *a priori* estimate of the spacecraft-feature relative pose along with a geometric camera model[8] is used. This allows the creation of correspondence between the topography pixels and the image pixels, which is used to map the intensity values to the surface (using bilinear interpolation).

The extraction process can lead to errors when topography pixels that should not be visible to the camera are assigned an intensity value due to errors in the *a priori* pose and noise in the image. Therefore, SPC Autoregister uses a filter to remove these pixels from consideration in the correlation. There are two primary checks in this filter. The first is to ignore pixels where the reflection or incident angles are greater than 90 degrees. The second check is for occlusion and shadow, which performs something analogous to a ray cast to the camera and the sun, checking for occlusion by other topography in the feature. SPC Autoregister does not check for shadows cast by other topography outside the feature.

Again we note that the resulting extracted image is locked to the feature topography, which is where the normalized cross correlation is performed.

Background of GIANT SFN

Similar to SPC Autoregister, GIANT SFN identifies bearing measurements to known surface features using normalized cross correlation. In contrast, however, GIANT SFN uses a single bounce ray trace to render the predicted template in the image space, where normalized cross correlation is then performed.

Template Rendering using Ray Tracing The rendering of the predicted templates for each feature takes place wholly in the image frame via a single bounce ray trace. Rays for each image pixel are traced through a model of the camera and into the scene where the closest topography intersection is identified. For each ray that intersects the topography, a bounce towards the light source (in this case the sun) is performed, checking for further intersections. When the reflected ray experiences

a second intersection, then the ray is labeled as being shadowed; whereas, if no second intersection occurs, the ray is considered illuminated. This process is shown in Figure 3.

Notably, this process produces a template in the image space.

Registration   Once the rendered template is generated, normalized cross correlation is performed with the image.[6] The typical normalized cross correlation algorithm is employed, with the exception that 2 masks are used to ignore pixels that are not part of the rendered feature topography. The first mask includes pixels where a ray through the pixel intersected the topography (this rejects pixels in the template that are outside the feature topography). The second mask includes pixels in the image which are of empty space. In the correlation process, the logical OR of these two masks is taken to determine which pixels should be included in the computation of the Pearson product coefficient. An example of this process is shown in Figure 4.

Comparison Between the Techniques

While both of these techniques rely on registering a rendered template based on *a priori* information with the image, there are a few significant differences in the details. These are organized into Table 1 to provide an easy comparison between the two methods.

EXPERIMENTAL COMPARISON OF THE METHODS

In this section we describe the analysis we performed to evaluate the relative strengths of the two tools. We begin by describing our general analysis steps, and then provide a discussion of the shape models and images used in that analysis.

Analysis Methodology

In order to compare the performance of the SPC Autoregister and GIANT SFN algorithms for extracting observables of surface features from images, we applied the tools on images from the OSIRIS-REx mission to asteroid Bennu using multiple sets of surface features developed in different ways. For each image and each set of surface features, we used each tool to extract observations of the visible features in the image. We allowed each tool to perform a single PnP solution to correct *a priori* errors before extracting the final observables.

We set tuning parameters on each tool based on how they were used during operations, with the important parameters described in Tables 2 and 3. In general, we tried to tune both algorithms similarly; however, in some instances we had to deviate slightly, generally due to minor differences in definitions between the tools.

With the observables and PnP solutions generated from each tool, we then analyzed the results. The first step was to reject outliers from each dataset (by rejecting all observations outside 3 sigmas for each image using median absolute deviation). With outliers rejected, we could then turn to examining the results.

We first examined the observed minus computed residuals using the PnP solution for the computed values. As part of this analysis, we plotted the residual magnitude in multiple meaningful directions (including the sun direction, the row direction, the column direction, and the image distortion gradient direction). This enabled us to see how consistent observations are in a single image, since the PnP solutions should correct any errors that can be corrected by a rigid transformation.

Next, we considered the distribution of residual statistics (including the median residual per-image and the standard deviation of the residuals per-image) for each image. This provides another check on the consistency of the observations in single images. Similarly, we examined the distribution of the correlation scores in each image, and the distribution of the number of features identified in each image. These metrics provide insight into how well we are matching the features to the image and if one tool is finding significantly more features than the other.

Finally, we looked at a direct difference. Here we identified all features that were found by both tools in the same image and differenced the observed locations, plotting the results in multiple frames and as function of multiple parameters. This analysis shows us any differences between the observed location for the feature in the image between the two tools and highlights any biases or trends.

In total, these analyses provide insight into the relative performance of the tools, as well as their overall performance in a range of image conditions and for a number of different feature sets.

Description of Data

For this analysis we used NavCam 1 images[9] of asteroid Bennu captured during the Orbital A, Detailed Survey, and Orbital B OSIRIS-REx mission phases. Orbital A was a 1.82 km frozen terminator orbit which ranged from about 1.5 to 2.1 km from the center of mass of Bennu.[10,11] Detailed Survey consisted of a number of 3–5 km flybys around Bennu, spanning solar phase angles from near 0 to 100 degrees.[11,12] Orbital B was a 0.93 km near circular, near-terminator orbit.[11] In all, these images provide a wide range of viewing conditions with which to test the tools.

In addition to the imaging set, we used 3 feature sets of Bennu generated using 3 different techniques. The first shape model we used, referred to as SPC v42, was the official operational shape model used by the navigation team beginning in Orbital B through the rest of proximity operations.[13] This set was built entirely with SPC surface modelling capabilities applied to images of Bennu and included 14,812 features at 5 cm, 14 cm, 35 cm, 75 cm, and 120 cm ground sample distances with the distribution shown in Figure 5. All features had relative albedo values as estimated through the SPC processes.

The next shape model we used, referred to as OLA v21, was a shape model generated using only data from the OSIRIS-REx laser altimeter (OLA).[13–15] This set included 36,450 features at 5 cm, 15 cm, 35 cm, 75 cm, and 150 cm ground sample distances (tiled from a global 5 cm shape model) with the distribution shown in Figure 5. These features include topography only (no relative albedo values).

The final shape model we used, referred to as SPCOLA, was a shape model generated using both images with SPC and OLA data. This feature set was built by adding the OLA-only features to the SPC-only features and then iterating in SPC to create consistency and produce a final model[13] and included 147,255 features at 0.4 cm, 0.8 cm, 2.0 cm, 5 cm, 12 cm, 20 cm, 30 cm, 75 cm, and 150 cm ground sample distances with the distribution shown in Figure 5. These features included both topography and relative albedo values.

In general, these shape models provide feature sets that are representative of those typically available to the tools, providing a useful sample for analysis. The SPC v42 model provides good topography and relative albedo, though it tends to slightly underestimate boulders and craters as photoclinometric approaches tend to do, which can sometimes cause problems with high incident angles. The OLA v21 model provides excellent topography, but doesn't provide relative albedo values, which

can be critical, especially in low-phase situations. The SPCOLA model theoretically provides the best of both the previous models, but takes the longest and the most amount of data to generate, which may be infeasible for some missions. This also is still an iterative approach that does not directly fuse the two data types together when constructing the features.

RESULTS

We now present the experimental results of applying the methods to the OSIRIS-REx images of Bennu. We examine both the results reported by the tools and the raw observations output by the tools to be fed to orbit determination software. To begin, we examine the per-image statistics from the observed minus computed residuals from the PnP solution for each image and the statistics on the correlation scores between the templates and images. We then examine the distribution of the residuals across all images to see if the observations are self-consistent with a rigid transformation. Finally, we difference the observed locations for common features found in the same images by both tools to see if any systematic differences result, since the observed locations are what are fed to the OD software.

Image Level Statistics

GIANT SFN with OLA V21 The per-image statistics of applying GIANT SFN to the OSIRIS-REx images with the OLA V21 shape model are shown in Figures 6 and 7. The residuals are generally unbiased (after the PnP solution) with per-image standard deviations less than 1 pixel, which is in line with expectations. In addition, the correlation scores between the templates and the images are generally around 0.8, with small standard deviations per-image, showing good agreement between the templates and the images. The number of features identified per-image is usually between 20 and 100, with some images stretching upwards to 350. In all, this shows good per-image performance for the GIANT SFN algorithm with the OLA V21 shape model.

SPC Autoregister with OLA V21 The per-image statistics of applying SPC Autoregister to the OSIRIS-REx images with the OLA V21 shape model are shown in Figures 8 and 9.

There appears to be a slight (but mostly negligible) bias in the row direction (after the PnP solution) with per-image standard deviations less than 1 pixel, which is in line with expectations. In addition, the correlation scores between the templates and the images are generally around 0.45, with small standard deviations per-image, showing passable agreement between the templates and the images. The lower correlation scores are primarily because SPC does not consider shadows in the correlation and the SPC V42 data did not include albedo information, which is much more important in cases where shadows are not included. The number of features identified per-image is usually between 20 and 100, with some images stretching upwards to 350, very similar to the results from GIANT SFN. In all this shows fairly good per-image performance for the SPC Autoregister algorithm with the SPC V42 shape model that is comparable to the GIANT SFN performance with the same model.

GIANT SFN with SPC V42 The per-image statistics of applying GIANT SFN to the OSIRIS-REx images with the SPC V42 shape model are shown in Figures 10 and 11. The residuals are generally unbiased (after the PnP solution) with per-image standard deviations less than 1 pixel, which is in line with expectations. We do note that the standard deviations are slightly higher for this shape model than for the OLA V21 model with GIANT SFN, possibly due to the greater number of features included in this model. In addition, the correlation scores between the templates

and the images are generally around 0.7, with small standard deviations per-image, showing good agreement between the templates and the images. The number of features identified per-image is usually between 20 and 200, with some images stretching upwards to 1000. We note that the increase in the number of features identified is due to the SPC V42 model having many more features with ground sample distance greater than 5 cm than the OLA V21 model. In all this shows good per-image performance for the GIANT SFN algorithm with the SPC V42 shape model.

SPC Autoregister with SPC V42 The per-image statistics of applying SPC Autoregister to the OSIRIS-REx images with the SPC V42 shape model are shown in Figures 12 and 13. There appears to be a slight (but mostly negligible) bias in the row direction (after the PnP solution) with per-image standard deviations less than 1 pixel, which is in line with expectations. In addition, the correlation scores between the templates and the images are generally around 0.7, with small standard deviations per-image, showing good agreement between the templates and the images. The number of features identified per-image is usually between 20 and 1000, with some images stretching upwards to 5000. We note that this is significantly more features found per-image than GIANT SFN with the same shape model, indicating that the GIANT SFN filter could be made more permissive. That being said, both techniques generate plenty of observations per-image for use in navigation. In all, this shows good per-image performance for the SPC Autoregister algorithm with the SPC V42 shape model that is comparable to the GIANT SFN performance with the same model.

GIANT SFN with SPCOLA The per-image statistics of applying GIANT SFN to the OSIRIS-REx images with the SPCOLA shape model are shown in Figures 14 and 15. The residuals are generally unbiased (after the PnP solution) with per-image standard deviations less than 1 pixel, which is in line with expectations. In addition, the correlation scores between the templates and the images are generally around 0.7, with small standard deviations per-image, showing good agreement between the templates and the images. The number of features identified per-image is usually between 20 and 250, with some images stretching upwards to 1750. We note that the increase in the number of features identified is due to the SPCOLA model having many more features with ground sample distance greater than 5 cm than the OLA V21 model and the SPC V42 model. In all this shows good per-image performance for the GIANT SFN algorithm with the SPCOLA shape model.

SPC Autoregister with SPCOLA The per-image statistics of applying SPC Autoregister to the OSIRIS-REx images with the SPCOLA shape model are shown in Figures 16 and 17. There appears to be a slight (but mostly negligible) bias in the row direction (after the PnP solution) with per-image standard deviations less than 1 pixel, which is in line with expectations. In addition, the correlation scores between the templates and the images are generally around 0.7, with small standard deviations per-image, showing good agreement between the templates and the images. The number of features identified per-image is usually between 20 and 2000, with some images stretching upwards to 10000. We again note that this is significantly more features found per-image than GIANT SFN with the same shape model, indicating that the GIANT SFN filter could be made more permissive. That being said, both techniques continue to generate plenty of observations per-image for use in navigation. In all, this shows good per-image performance for the SPC Autoregister algorithm with the SPCOLA shape model that is comparable to the GIANT SFN performance with the same model.

PnP Solution Residuals

In this section we show the residuals after the PnP solution by each tool in a number of frames to identify patterns.

GIANT SFN with OLA V21   The post-PnP residuals from GIANT SFN for all images using the OLA V21 Shape model are shown in Figure 18. These plots utilize hex bins, where colored bins are used to show the number of points for each region of the plot with deeper colors indicating more points (colored using a logarithmic scale). Where appropriate, we superimpose a rolling median in gray and the +/- 1 sigma bounds about the median in red to highlight trends. In Figure 18 (and in future similar figures), each row represents a different direction the residuals are expressed in, as labeled on the left, and each column represents a different independent variable we are plotting the residuals against, labeled on the bottom with units (i.e., the top right plot shows the sun direction residuals as a function of the camera distortion magnitude at the identified feature location). In addition, in the plot, the "Radial Direction" is the direction from the observed location to the predicted Bennu center of mass in the image. The "Distortion Direction" is the direction that the camera distorts pixels at the observed location, and the "Distortion Magnitude" is the magnitude of the distortion field at this point.

In general, the residuals are very small with no major trends. In fact, in all directions, the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.2 pixels. Where trends exist, the trend median is normally within 1 sigma of 0, and at worst within 2 sigma of 0.

The strongest (but still small) subpixel trend is in the camera distortion and radial directions as a function of the distortion magnitude, which we believe may be caused by rendering the templates in the wrong region of the camera distortion field due to errors in the *a priori* pose knowledge (note that the radial direction and the distortion direction frequently are similar, especially in images where Bennu is near the center of the field of view of the camera, as it is in the majority of the dataset). There are also small subpixel trends in the sun direction and center of mass radial direction (which are also commonly similar for many images due to the terminator orbits and the fact that the sun was usually fixed to come from the same direction in the images due to the spacecraft attitude). These trends (especially the sun direction as a function of the incident angle) may be caused by the inclusion of the shadows in the correlation, which may bias the results if the topography isn't exactly right. However, all of these trends are at the subpixel level and are generally only pronounced in the extremes, where data is very sparse to begin with, and therefore are unlikely to have major contributions to most navigation analysis.

SPC Autoregister with OLA V21   The post-PnP residuals from SPC Autoregister for all images using the OLA V21 Shape model are shown in Figure 19. In general, the residuals are very small with no major trends. As with the GIANT SFN results, in all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.2 pixels. Where trends exist, the trend median is within 1 sigma of 0.

There is no longer a strong trend in the distortion direction, which may be due to SPC Autoregister correlating on the feature (where there is no image distortion) as opposed to in the image. There is a small subpixel trend in the sun direction as a function of the incidence angle. Interestingly, the sun direction trend is most pronounced where the most data is available, unlike in the GIANT SFN results where the trend is strongest where there is the least amount of data. These trends may be caused by the lack of albedo information for the features, which SPC Autoregister relies upon heavily, though further in depth analysis is required to truly understand. We also note that residuals with very high phase angles (¿90 degrees) perform much worse than in the GIANT SFN results. This is likely due to high phase angles resulting in elongated shadows and minimal unshadowed regions to correlate with. However, all of these trends are at the small subpixel level and therefore

are unlikely to have major contributions to most navigation analysis.

GIANT SFN with SPC V42   The post-PnP residuals from GIANT SFN for all images using the SPC V42 Shape model are shown in Figure 20. In general, the residuals are very small with no major trends. In all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.4 pixels. Where trends exist, the trend median is within 1 sigma of 0.

As with the OLA V21 results, the strongest trend is still in the distortion direction as a function of the distortion magnitude, though it is smaller for this shape model. Beyond that there is a very small trend in the sun direction as a function of the incidence angle, but not much else. As before, the sun direction trend may be caused by the inclusion of the shadows in the correlation, which may bias the results if the topography isn't exactly right, though interestingly here the trend is opposite that seen in the OLA V21 model.

SPC Autoregister with SPC V42   The post-PnP residuals from SPC Autoregister for all images using the SPC V42 Shape model are shown in Figure 21. In general, the residuals are very small with no major trends. As with the GIANT SFN results, in all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.35 pixels. Where trends exist, the trend median is within 1 sigma of 0.

There is only one small trend of note, in the sun direction as a function of the incident angle, and this has flattened from the OLA V21 model. The worst performance is still at high phase angles; however, it is much more in line with the SFN performance, indicating that the albedo information is critical to SPC. This is also indicated by the best performance being at the low phase angles.

GIANT SFN with SPCOLA   The post-PnP residuals from GIANT SFN for all images using the SPCOLA Shape model are shown in Figure 22. In general, the residuals are very small with no major trends. In all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.2 pixels. Where trends exist, the trend median is within 1 sigma of 0.

There is no notable trend in the distortion direction. This may be due to the substantially greater number of features, allowing the PnP solution to better correct errors in the *a priori* pose knowledge for the second pass. There are also small subpixel trends in the sun direction and center of mass radial direction as a function of the incidence angle (which are very similar for many images due to the terminator orbits). These trends (especially the sun direction as a function of the incident angle) may be caused by the inclusion of the shadows in the correlation, which may bias the results if the topography isn't exactly right. All of these trends are at the subpixel level and are generally only pronounced in the extremes, where data is very sparse to begin with, and therefore are unlikely to have major contributions to most navigation analysis.

SPC Autoregister with SPCOLA   The post-PnP residuals from SPC Autoregister for all images using the SPCOLA Shape model are shown in Figure 23. In general, the residuals are very small with no major trends. In all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.2 pixels.

There are no clearly observable trends in this dataset. This could indicate that the large number of features for the model enabled the PnP solution to better correct for errors in the *a priori* pose knowledge, or it could indicate that the data is best suited for this technique.

9

Differenced Observables

In this section we show the difference between the observed feature location for common features found by each tool in each image. For each identified feature location pair, we difference the SPC Autoregister value from the GIANT SFN value before plotting in different frames as before to examining trends.

Differenced OLA V21   The difference for common features found in each image by the two tools using the OLA V21 Shape model are shown in Figure 24.

In general, the residuals are small with only minor trends. In all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.2 pixels. Where trends exist, the trend median is within 1 sigma of 0.

The trends in the distortion direction and the center of mass direction seen in the PnP solution residuals for each tool have largely flattened out, indicating the trends are most pronounced in features observed by only one of the tools. The trend in the sun direction is exaggerated from the PnP solution trends. This indicates that the PnP solution trends either are strongest in the commonly observed features or are in opposite directions. This trend is likely due to the difference in the way that SPC Autoregister and GIANT handle shadows. However, even with the exaggerated trend here, in the extreme it is less than half a pixel, with the vast majority of data being much closer to zero.

Differenced SPC V42   The difference for common features found in each image by the two tools using the SPC V42 shape model are shown in Figure 25.

In general, the residuals are small with only minor trends. In all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.2 pixels. Where trends exist, the trend median is within 1 sigma of 0.

The trends in the distortion direction and the center of mass direction seen in the PnP solution residuals for each tool are largely gone, indicating the trends are most pronounced in features observed by only one of the tools. The trend in the sun direction is somewhat exaggerated from the PnP solution trends but also more centered near 0. This trend is likely due to the difference in the way that SPC Autoregister and GIANT handle shadows. However, even with the exaggerated trend here, the most extreme median residuals are less than a quarter pixel, with the vast majority of data being much closer to zero.

Differenced SPCOLA   The difference for common features found in each image by the two tools using the SPCOLA shape model are shown in Figure 26. In general, the residuals are small with only minor trends. In all directions the mean and median are less than 0.1 pixels in absolute value and the standard deviation is around 0.3 pixels. Where trends exist, the trend median is within 1 sigma of 0.

The trends in the distortion direction and the center of mass direction seen in the PnP solution residuals for each tool are largely gone, indicating the trends are most pronounced in features observed by only one of the tools. There is a small trend in the sun direction, similar to the one seen in the GIANT SFN PnP solution. This trend is likely due to the difference in the way that SPC Autoregister and GIANT handle shadows.

CONCLUSION

Overall both tools produce observations of surface features to subpixel accuracy and there is generally good agreement between the observations generated by the tools. Both tools show good consistency with the PnP solutions across all shape models. In general the results are unbiased for both tools and all shape models. This also extends to the differenced observables between the tools. There are a few trends, but they are generally at the subpixel level, within 1 sigma of 0, and are typically towards the extremes of the data where there are fewer observations allowing outliers to have larger sway.

Of the tool and shape model combinations analyzed, it appears that the SPC Autoregister and SPCOLA model are the most consistent, showing no major trends and having a small standard deviation. This is closely followed by the GIANT SFN and SPCOLA model, which has only very minor trends and a small standard deviation. That the SPCOLA model performs the best for both tools should come as no surprise since it should have the best topography and albedo combination. This illustrates the need for robust shape modelling capabilities that can fuse multiple data types.

For the other 2 shape models, the GIANT SFN tool performed slightly better with the OLA V21 model, which follows expectations since it heavily relies upon topography and shadowing, while the SPC Autoregister performed slightly better with the SPC V42 model, which again follows expectations since it more heavily relies upon albedo rather than shadowing. Based on this, in cases where you expect to have shape models with good topography but poor albedo information, you may prefer the GIANT SFN tool, whereas if you are building your models entirely from images, and are likely to have slightly smoothed topography because of that, you may prefer the SPC Autoregister tool.

All this being said, both tools and all shape models performed excellently and will produce observables that can be used to perform very accurate OD. However, there is still much analysis to be done to compare the tools. In this paper we relied on the PnP solution residuals, which demonstrate consistency with a rigid transformation, but this does not necessarily indicate consistency with a sequence of images taken from an orbit. We hope to explore this in the future by fitting this data in OD solutions to see if the residuals of the OD solution show similar trends. We also note that the PnP solutions (and for that matter the OD solutions) can be influenced if the overall shape model is not correct, resulting in bad feature locations. We are still investigating ways to determine which of these shape models most truly represent both the local and global topography of Bennu.

ACKNOWLEDGMENT

REFERENCES

[1] P. G. Antreasian, M. C. Moreau, C. D. Adam, A. French, J. Geeraert, K. M. Getzandanner, D. E. Highsmith, J. M. Leonard, E. J. Lessac-Chenen, A. H. Levine, *et al.*, "Early Navigation Performance of the OSIRIS-REx Approach to Bennu," *42nd AAS Rocky Mountain Section Guidance and Control Conference, 2019*, Univelt Inc., 2019, pp. 141–160.

[2] R. W. Gaskell, "Optical Navigation Near Small Bodies," *AAS/AIAA Spaceflight Mechanics Meeting, 2011*, Vol. 140, Univelt, Inc., 2011, pp. 1705–1717.

[3] A. Liounis, J. Swenson, J. Small, J. Lyzhoft, B. Ashman, K. Getzandanner, D. Highsmith, M. Moreau, C. Adam, P. Antreasian, *et al.*, "Independent Optical Navigation Processing for the OSIRIS-REx Mission Using the Goddard Image Analysis and Navigation Tool," *2nd RPI Space Imaging Workshop*, 2019.

[4] C. Gnam, A. Liounis, B. Ashman, K. Getzandanner, J. Lyzhoft, J. Small, D. Highsmith, C. Adam, J. Leonard, P. Antreasian, *et al.*, "A Novel Surface Feature Navigation Algorithm Using Ray Tracing," *2nd RPI Space Imaging Workshop*, 2019.

[5] A. J. Liounis, C. Gnam, J. Swenson, K. Getzandanner, J. Small, and J. Lyzhoft, "The Goddard Image Analysis and Navigation Tool," January 2022, https://github.com/nasa/giant.

[6] R. Brunelli, *Template Matching Techniques in Computer Vision: Theory and Practice*. John Wiley & Sons, 2009.

[7] A. S. McEwen, "Photometric Functions for Photoclinometry and Other Applications," *Icarus*, Vol. 92, No. 2, 1991, pp. 298–311.

[8] B. J. Bos, D. S. Nelson, J. Y. Pelgrift, A. J. Liounis, D. Doelling, C. D. Norman, R. D. Olds, C. W. May, R. Witherspoon, E. Church, D. Huish, C. D. Adam, E. Sahr, J. Kidd, K. Drozd, W. M. Owen, M. C. Moreau, L. T. Seals, J. Butt, D. LeDuc, K. Alkiek, L. R. Chevres-Fernandez, K. E. Gordon, K. Khlopenkov, C. Haney, R. Bhatt, J. Hikes, A. Wolfram, M. Caplinger, M. A. Ravine, K. Getzandanner, A. Bartels, and D. S. Lauretta, "In-Flight Calibration and Performance of the OSIRIS-REx Touch And Go Camera System (TAGCAMS)," *Space Science Reviews*, Vol. 216, No. 4, 2020, p. 71.

[9] B. Bos, M. Ravine, M. Caplinger, J. Schaffner, J. Ladewig, R. Olds, C. Norman, D. Huish, M. Hughes, S. Anderson, *et al.*, "Touch and Go Camera System (TAGCAMS) for the OSIRIS-REx asteroid sample return mission," *Space Science Reviews*, Vol. 214, No. 1, 2018, p. 37.

[10] D. R. Wibben, A. Levine, S. Rieger, J. V. McAdams, P. G. Antreasian, J. M. Leonard, M. C. Moreau, and D. S. Lauretta, "OSIRIS-REx frozen orbit design and flight experience," *AAS/AIAA Astrodynamics Specialist Conference*, Portland, ME, 2019, AAS 19-677.

[11] P. Antreasian, C. Adam, J. Leonard, E. Lessac-Chenen, A. Levine, J. McAdams, L. McCarthy, D. Nelson, B. Page, J. Pelgrift, E. Sahr, D. Wibben, B. Williams, K. Williams, K. Berry, K. Getzandanner, M. Moreau, S. Rieger, B. Ashman, D. Highsmith, and D. Lauretta, "OSIRIS-REx Proximity Operations and Navigation Performance at Bennu," *The 32nd AIAA/AAS Space Flight Mechanics Meeting*, San Diego, CA, 2022.

[12] D. Wibben, A. Levine, S. Rieger, J. McAdams, K. Getzandanner, P. Antreasian, J. Leonard, M. Moreau, and D. Lauretta, "Trajectory Design and Maneuver Performance of the OSIRIS-REx Detailed Survey of Bennu," *43rd Annual AAS Guidance and Control Conference*, 2020, pp. 20–151.

[13] O. Barnouin, M. Daly, E. Palmer, C. Johnson, R. Gaskell, M. Al Asad, E. Bierhaus, K. Craft, C. Ernst, R. Espiritu, H. Nair, G. Neumann, L. Nguyen, M. Nolan, E. Mazarico, M. Perry, L. Philpott, J. Roberts, R. Steele, J. Seabrook, H. Susorney, J. Weirich, and D. Lauretta, "Digital Terrain Mapping by the OSIRIS-REx Mission," *Planetary and Space Science*, Vol. 180, 2020, p. 104764, https://doi.org/10.1016/j.pss.2019.104764.

[14] M. G. Daly, O. S. Barnouin, C. Dickinson, J. Seabrook, C. L. Johnson, G. Cunningham, T. Haltigin, D. Gaudreau, C. Brunet, I. Aslam, A. Taylor, E. B. Bierhaus, W. Boynton, M. Nolan, and D. S. Lauretta, "The OSIRIS-REx Laser Altimeter (OLA) Investigation and Instrument," *Space Science Reviews*, Vol. 212, Oct 2017, pp. 899–924, 10.1007/s11214-017-0375-3.

[15] M. G. Daly, O. S. Barnouin, J. A. Seabrook, J. Roberts, C. Dickinson, K. J. Walsh, E. R. Jawin, E. E. Palmer, R. Gaskell, J. Weirich, T. Haltigin, D. Gaudreau, C. Brunet, G. Cunningham, P. Michel, Y. Zhang, R.-L. Ballouz, G. Neumann, M. E. Perry, L. Philpott, M. M. A. Asad, C. L. Johnson, C. D. Adam, J. M. Leonard, J. L. Geeraert, K. Getzandanner, M. C. Nolan, R. T. Daly, E. B. Bierhaus, E. Mazarico, B. Rozitis, A. J. Ryan, D. N. DellaGiustina, B. Rizk, H. C. M. Susorney, H. L. Enos, and D. S. Lauretta, "Hemispherical Differences in the Shape and Topography of Asteroid (101955) Bennu," *Science Advances*, Vol. 6, No. 41, 2020, p. eabd3649, 10.1126/sciadv.abd3649.
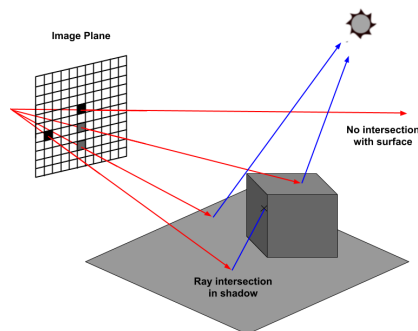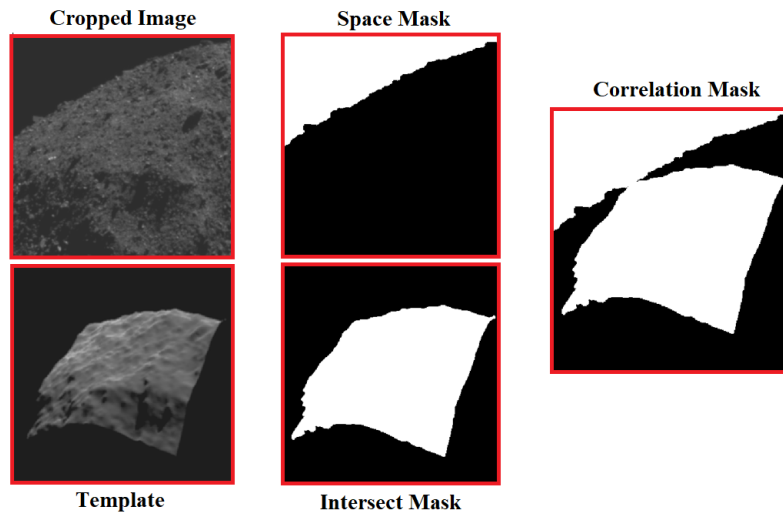
FIGURES
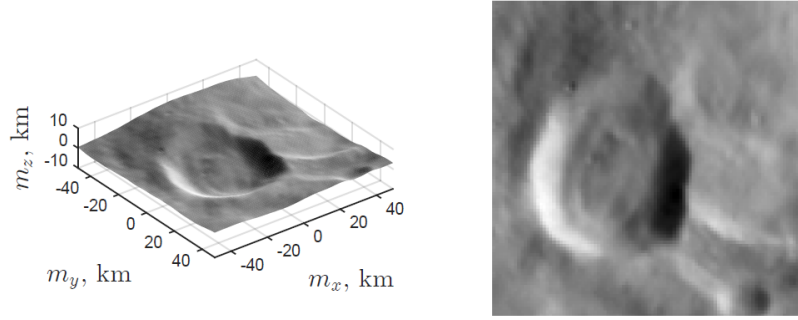
Figure 1. The predicated and extracted intensity values used in SPC navigation are on the actual surface of the feature being considered (left). This is also where the 2D cross correlation occurs using the top-down view of the landmark (right).
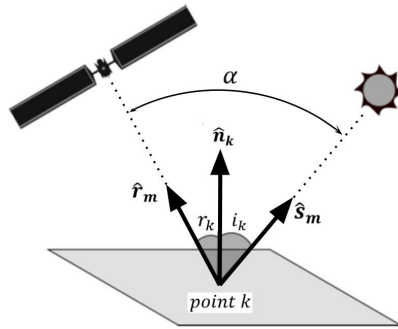

Figure 2. The geometry used in computing intensity from a BRDF. Reproduced with permission from Reference 4.


**Figure 3. The geometry of a single-bounce ray trace for a simple scene. Reproduced with permission from Reference 4.**

Cropped Image

Space Mask

Correlation Mask

Template

Intersect Mask

Figure 4. A diagram showing the generation of the correlation mask by determining a space mask from the cropped image and an intersection mask from the rendered template. The two masks are combined with a logical OR to determine which pixels to include in the computation of the correlation score for this template. We include both intersected pixels in the template and space pixels from the image to assist with correlating features that wrap beyond the visible limb in the image. Reproduced with permission from Reference 4.
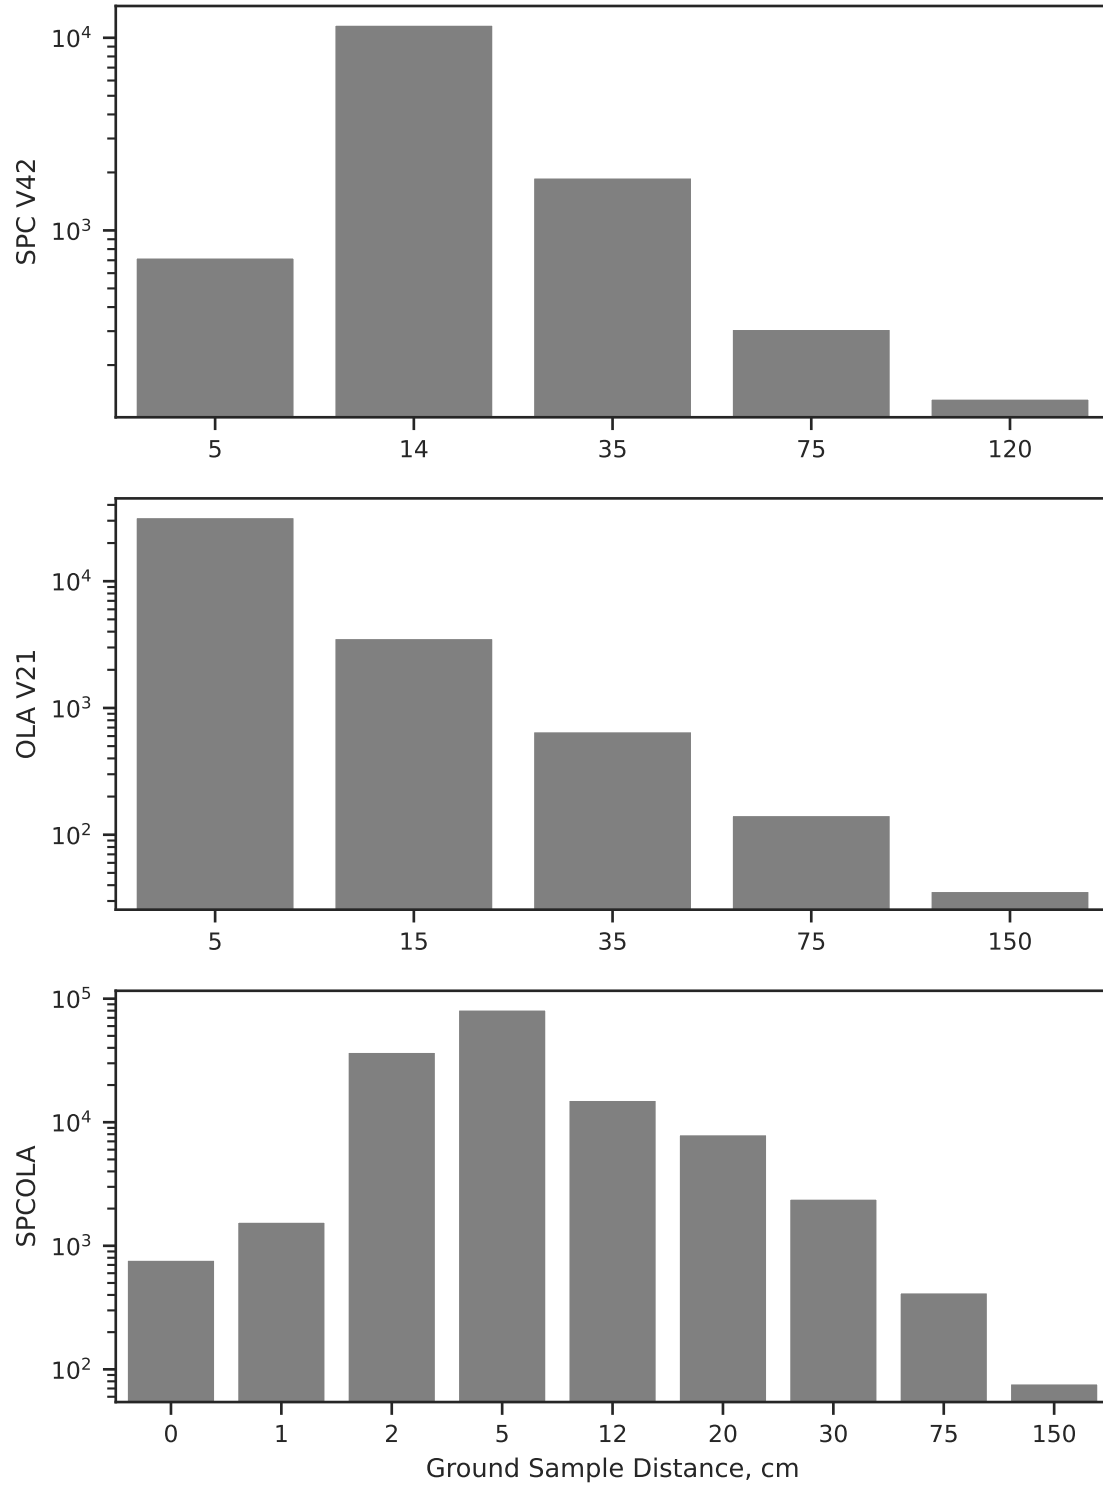
Figure 5. The number of features per ground sample distance in the used models.
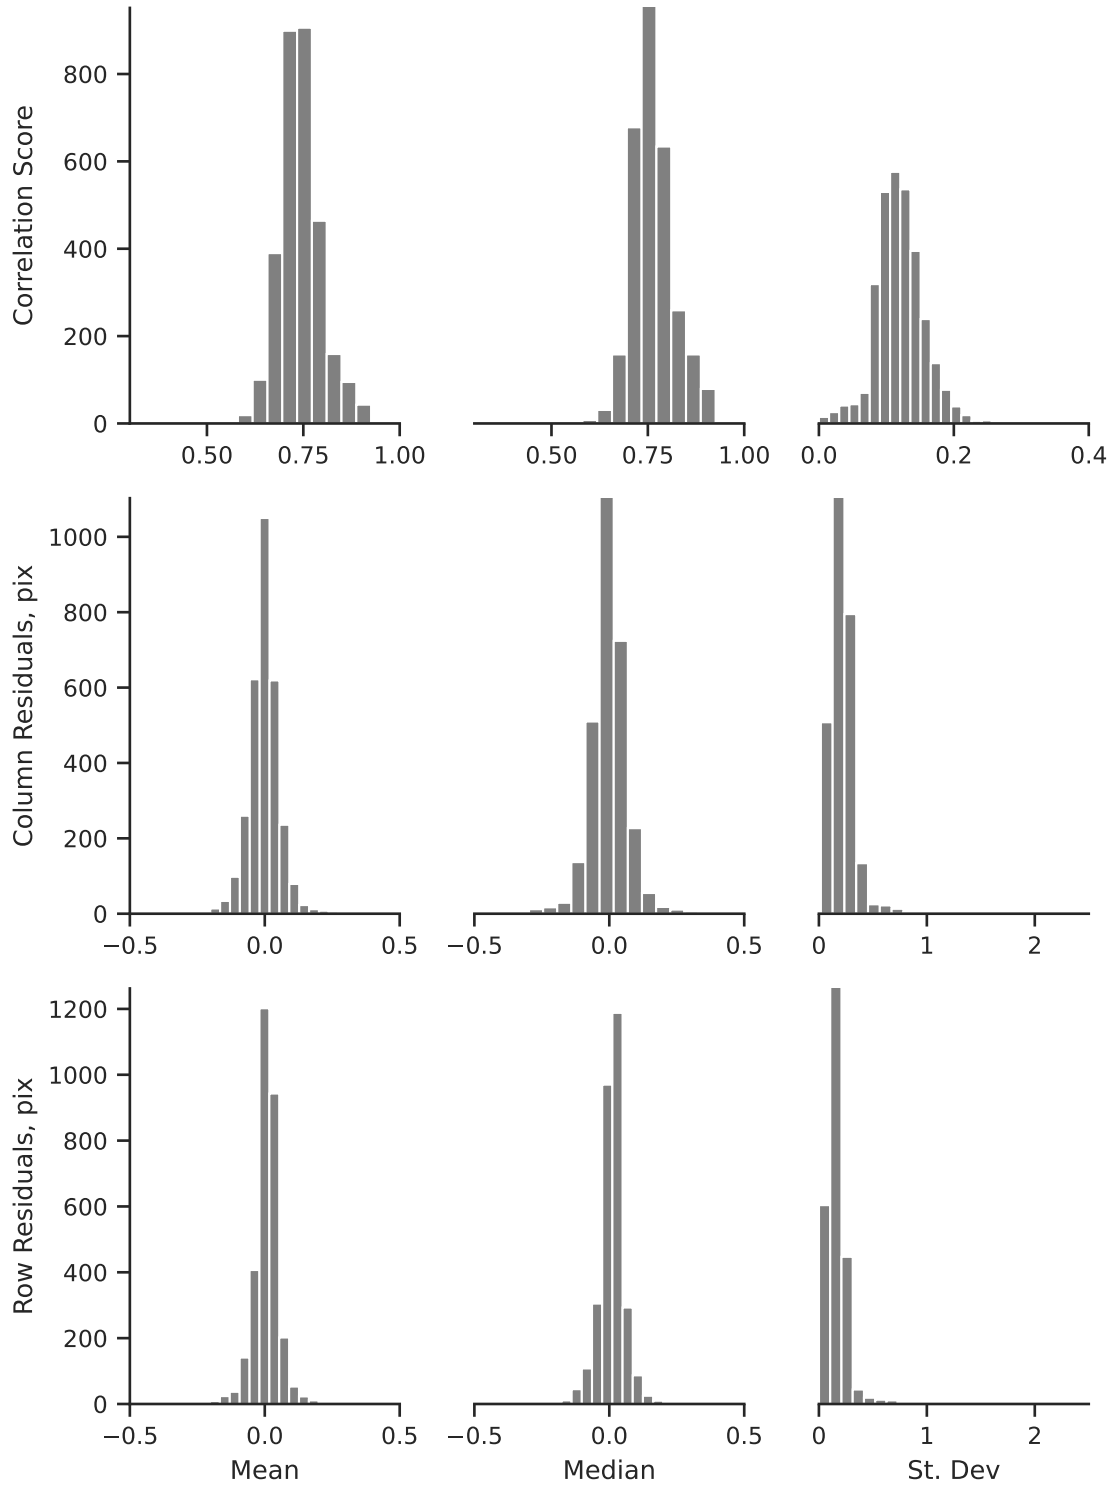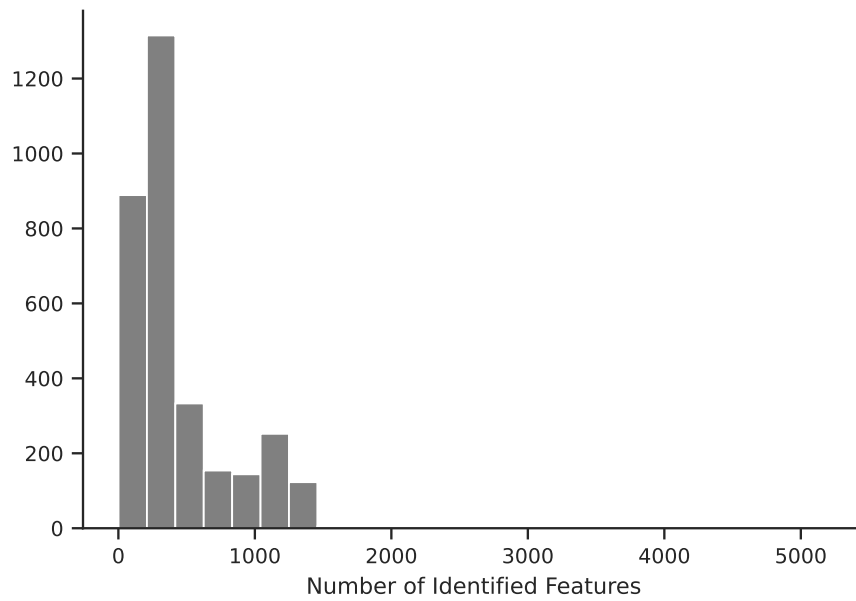
Figure 6. The per-image statistics from applying the GIANT SFN Algorithm with the OLA V21 shape model. Each row shows histograms of the residuals labeled on the left with each column showing the statistic labeled on the bottom (i.e., the center plot shows the per-image median of the column residuals in pixels).
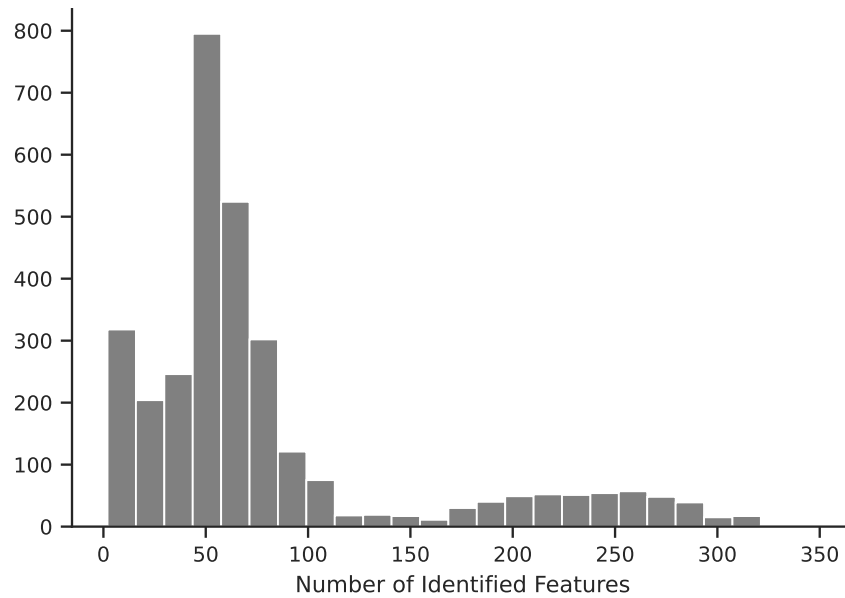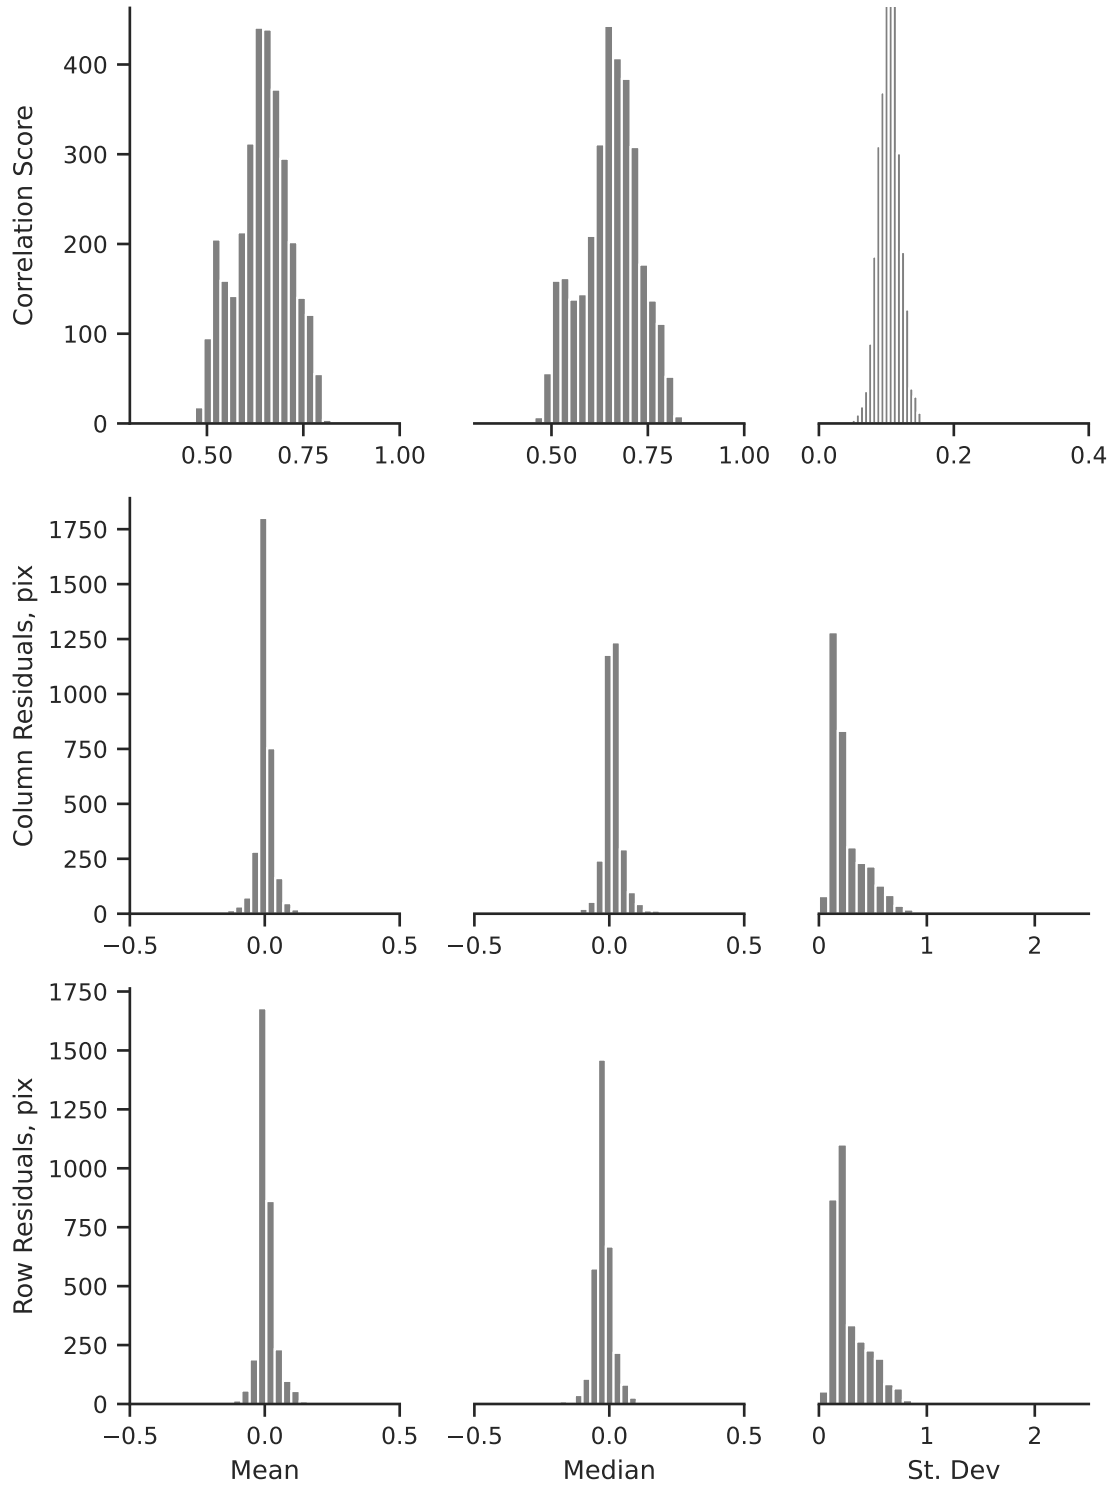
Figure 7.   The number of features identified per-image from applying the GIANT SFN algorithm with the OLA V21 shape model.
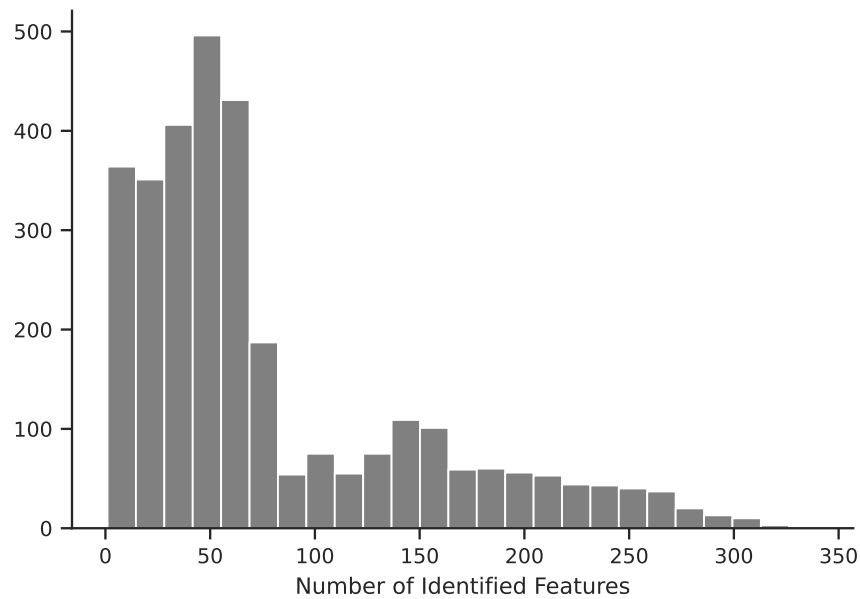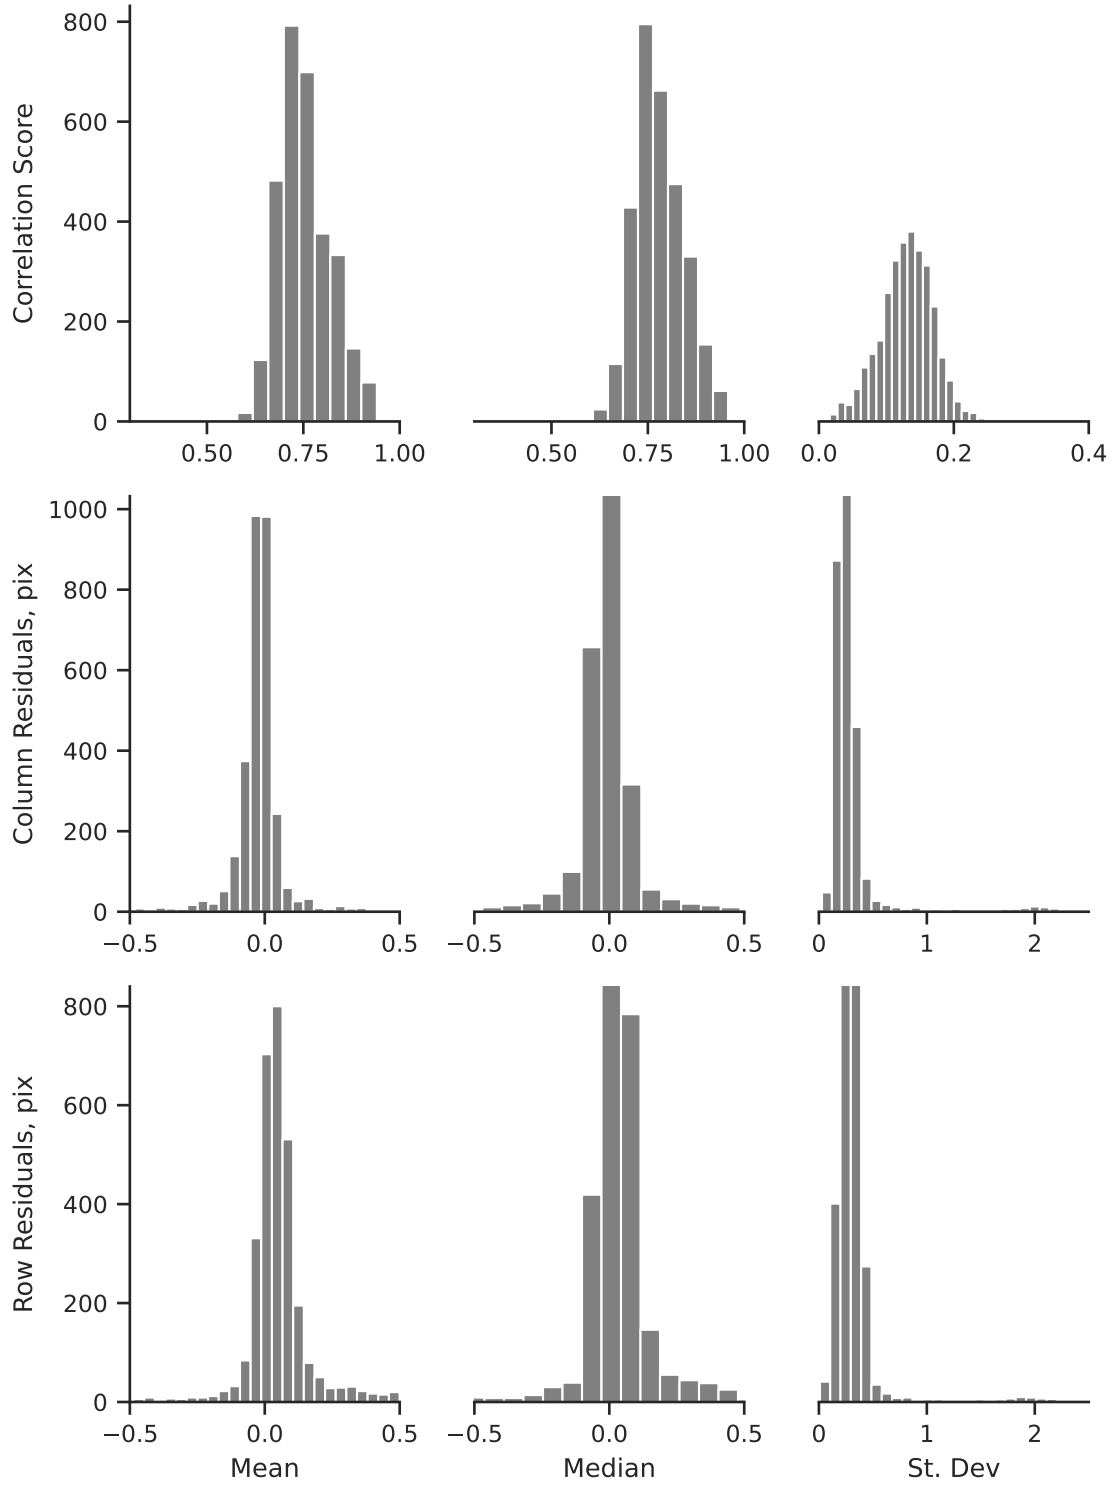

Figure 8.   The number of features identified per-image from applying the SPC Autoregister algorithm with the OLA V21 shape model.
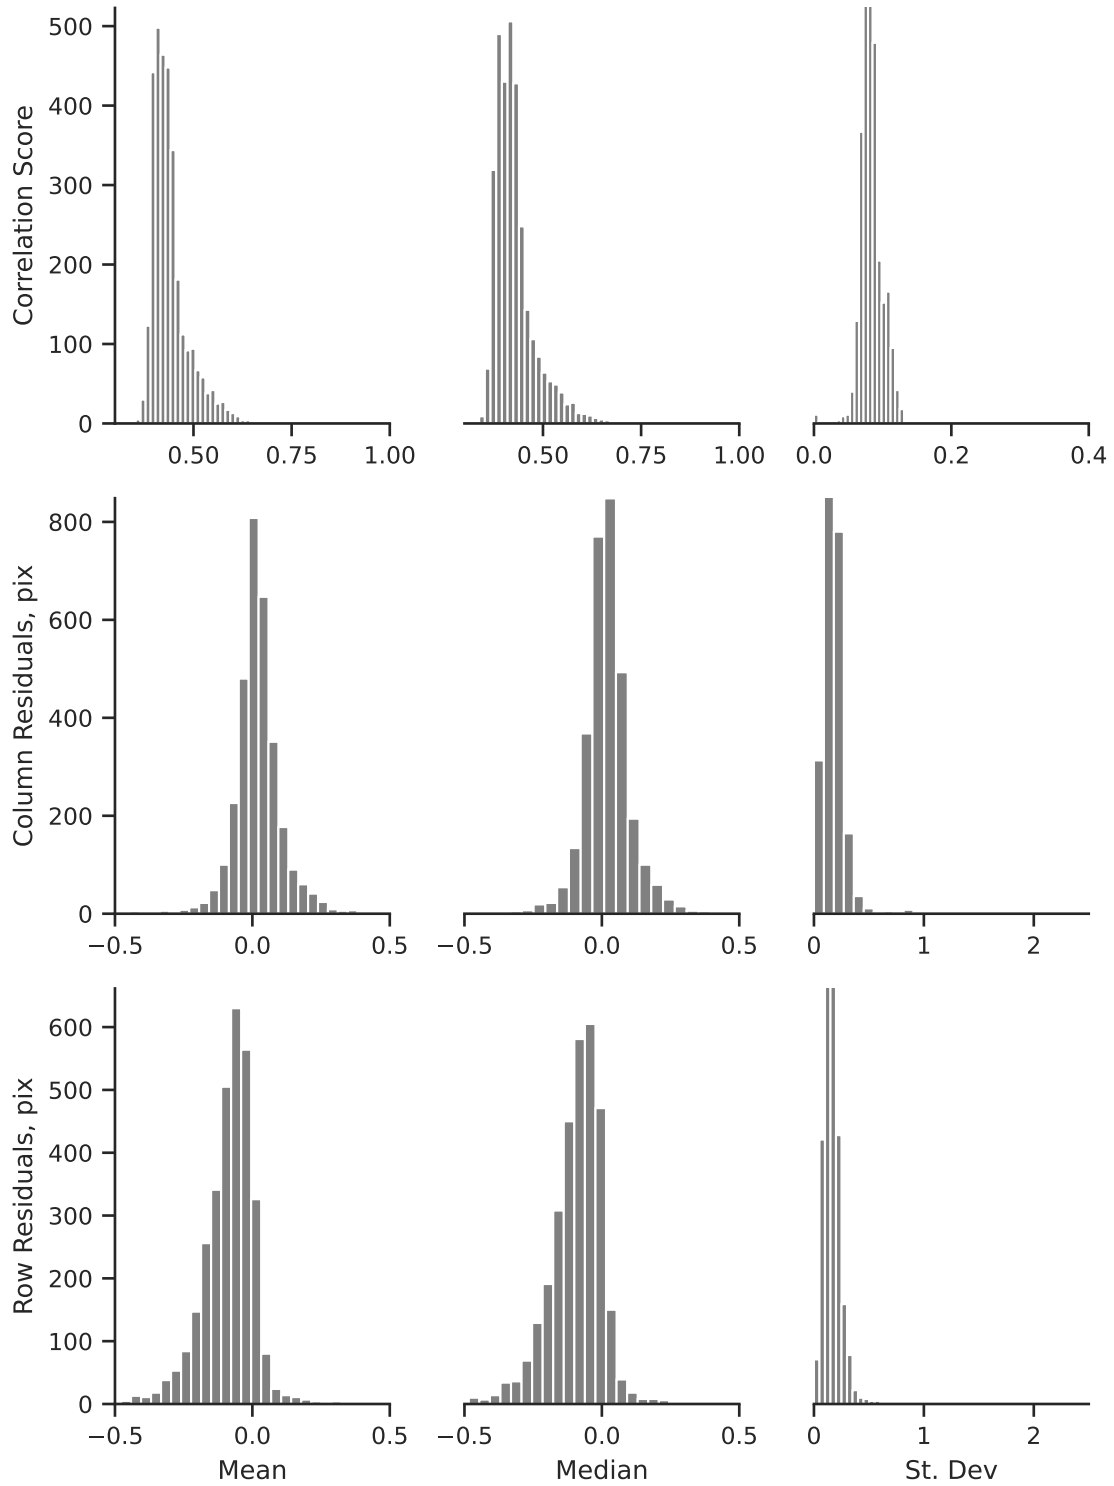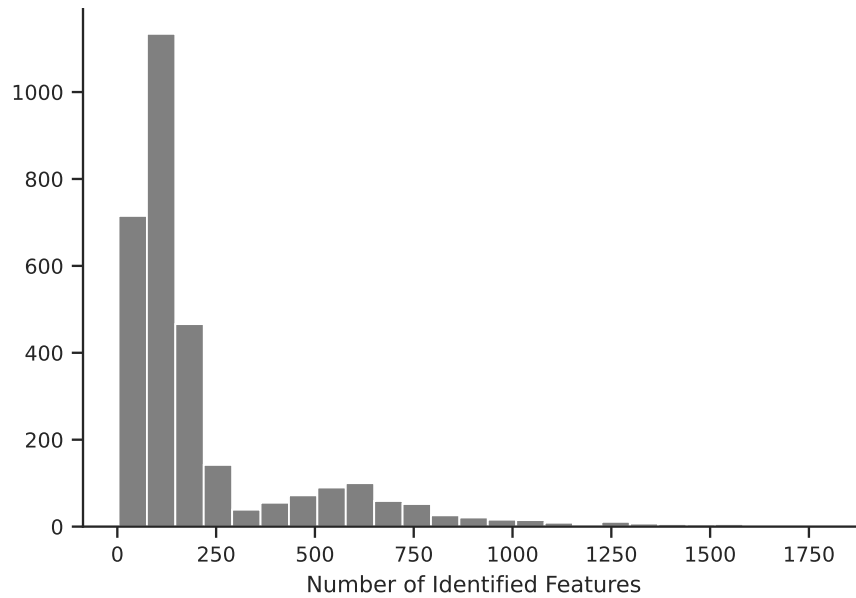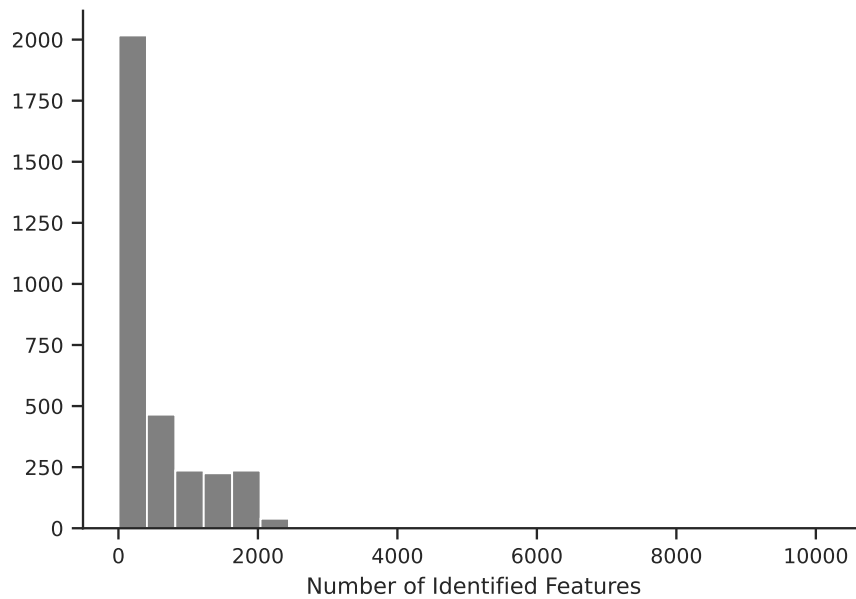
Figure 9. The per-image statistics from applying the SPC Autoregister algorithm with the OLA V21 shape model. Each row shows histograms of the residuals labeled on the left with each column showing the statistic labeled on the bottom (i.e., the center plot shows the per-image median of the column residuals in pixels).
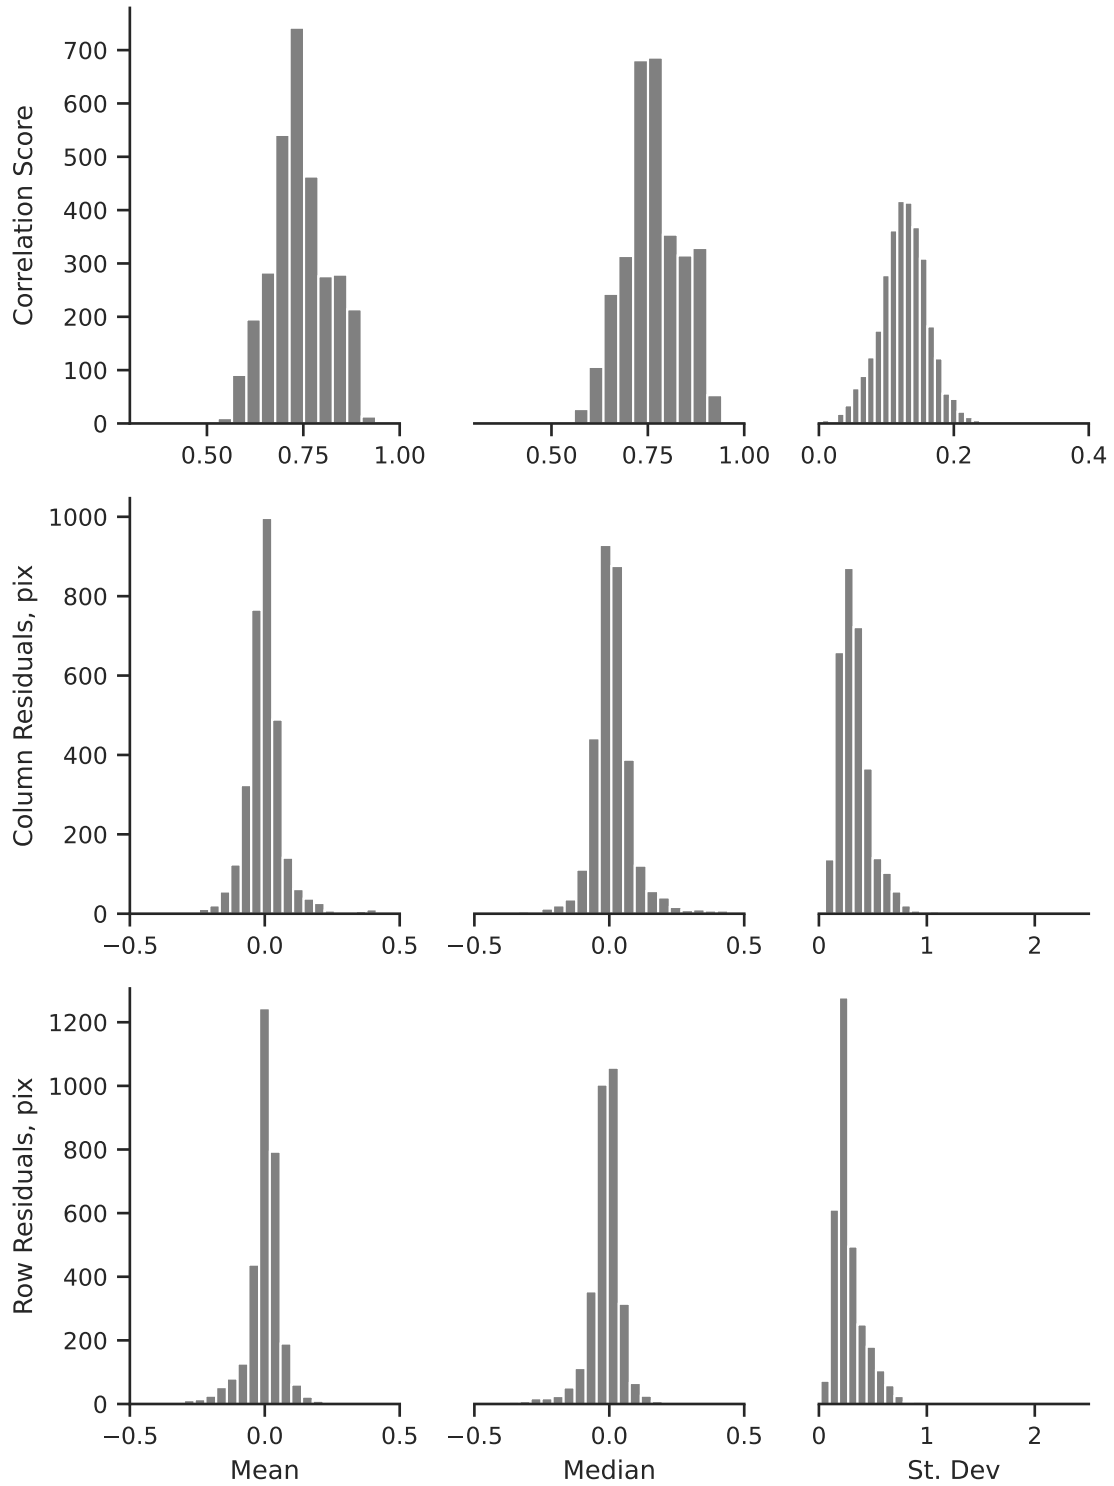
Figure 10. **The per-image statistics from applying the GIANT SFN Algorithm with the SPC V42 shape model. Each row shows histograms of the residuals labeled on the left with each column showing the statistic labeled on the bottom (i.e., the center plot shows the per-image median of the column residuals in pixels).**
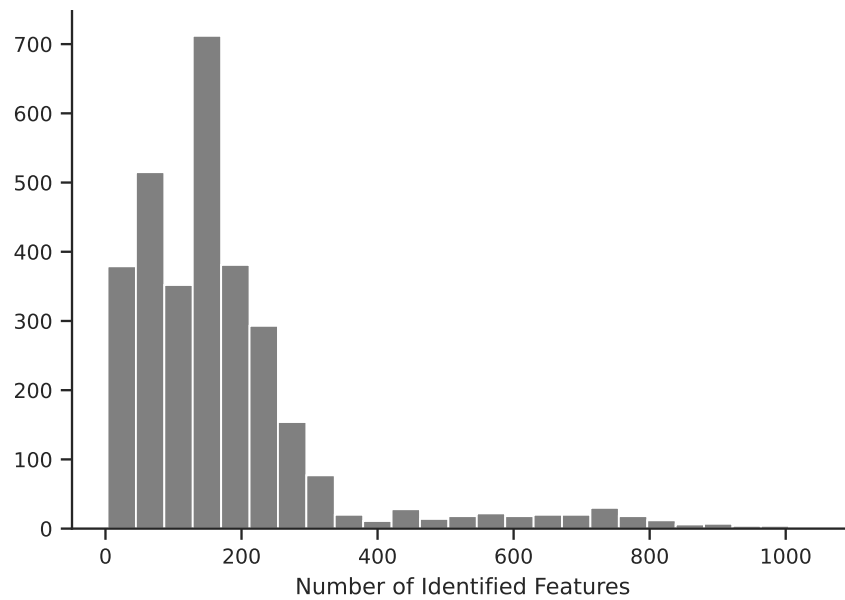
Figure 11. The number of features identified per-image from applying the GIANT SFN algorithm with the SPC V42 shape model.

Figure 12. The number of features identified per-image from applying the SPC Autoregister algorithm with the SPC V42 shape model.

Figure 13. **The per-image statistics from applying the SPC Autoregister algorithm with the SPC V42 shape model. Each row shows histograms of the residuals labeled on the left with each column showing the statistic labeled on the bottom (i.e., the center plot shows the per-image median of the column residuals in pixels).**

Figure 14.   The per-image statistics from applying the GIANT SFN Algorithm with the SPC V42 shape model. Each row shows histograms of the residuals labeled on the left with each column showing the statistic labeled on the bottom (i.e., the center plot shows the per-image median of the column residuals in pixels).

Figure 15.   The number of features identified per-image from applying the GIANT SFN algorithm with the SPCOLA shape model.


Figure 16.   The number of features identified per-image from applying the SPC Autoregister algorithm with the SPCOLA shape model.
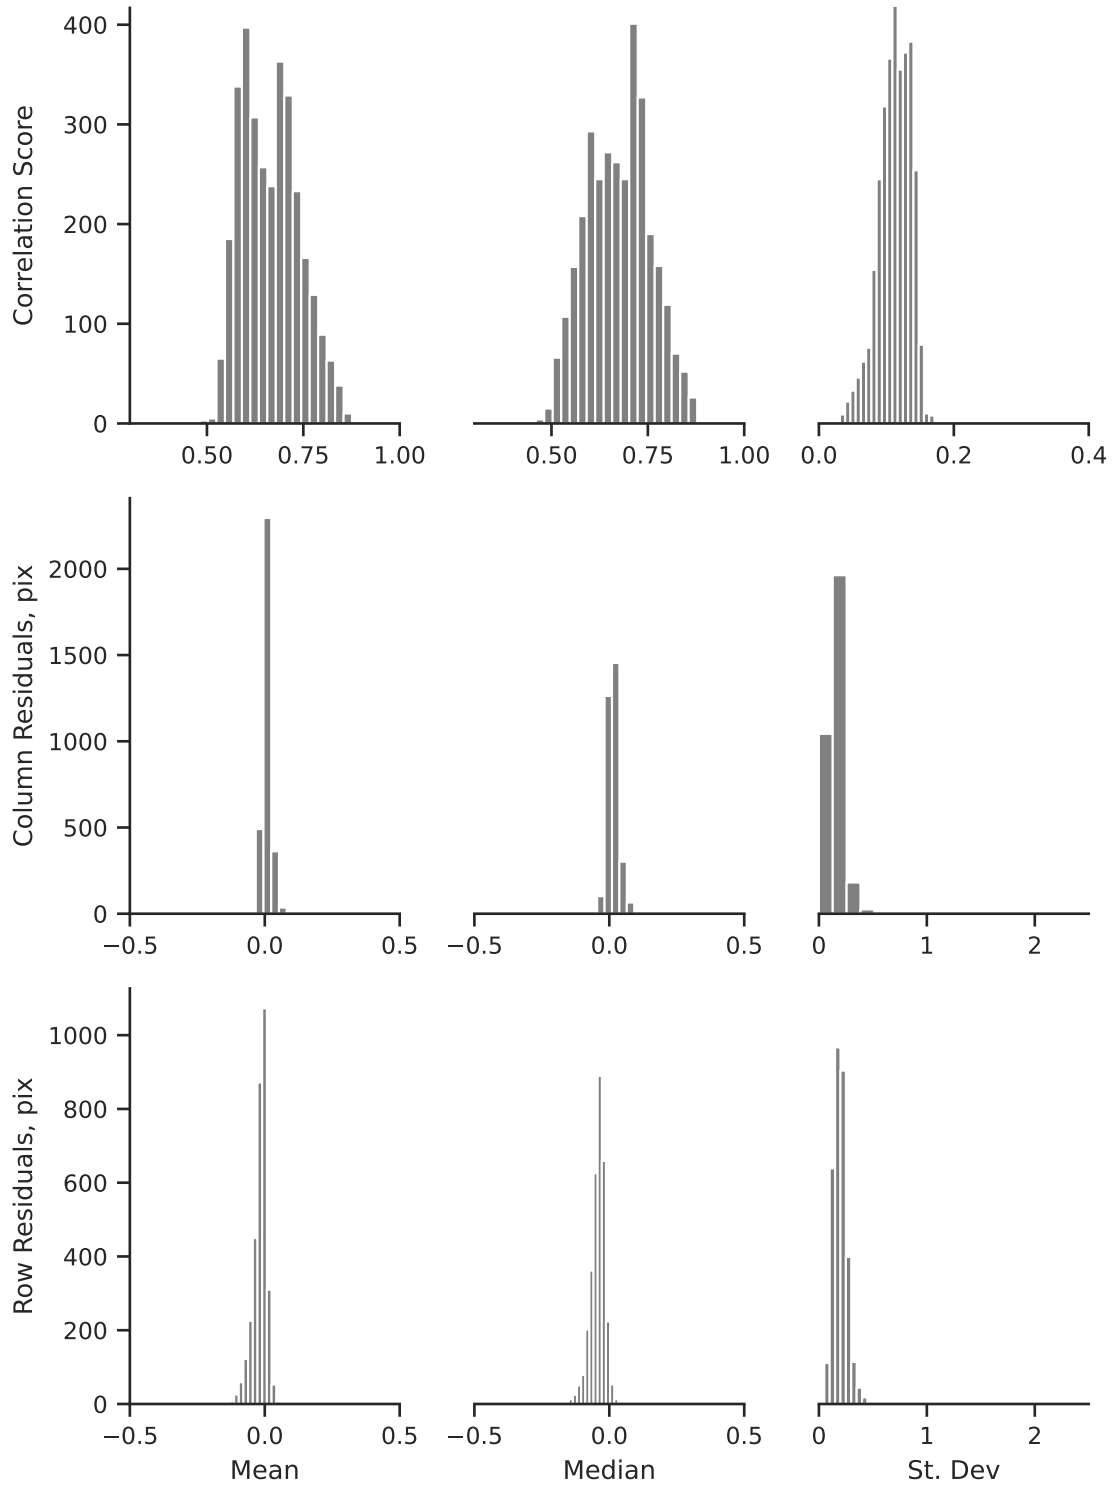
Figure 17. The per-image statistics from applying the SPC Autoregister algorithm with the SPCOLA shape model. Each row shows histograms of the residuals labeled on the left with each column showing the statistic labeled on the bottom (i.e., the center plot shows the per-image median of the column residuals in pixels).
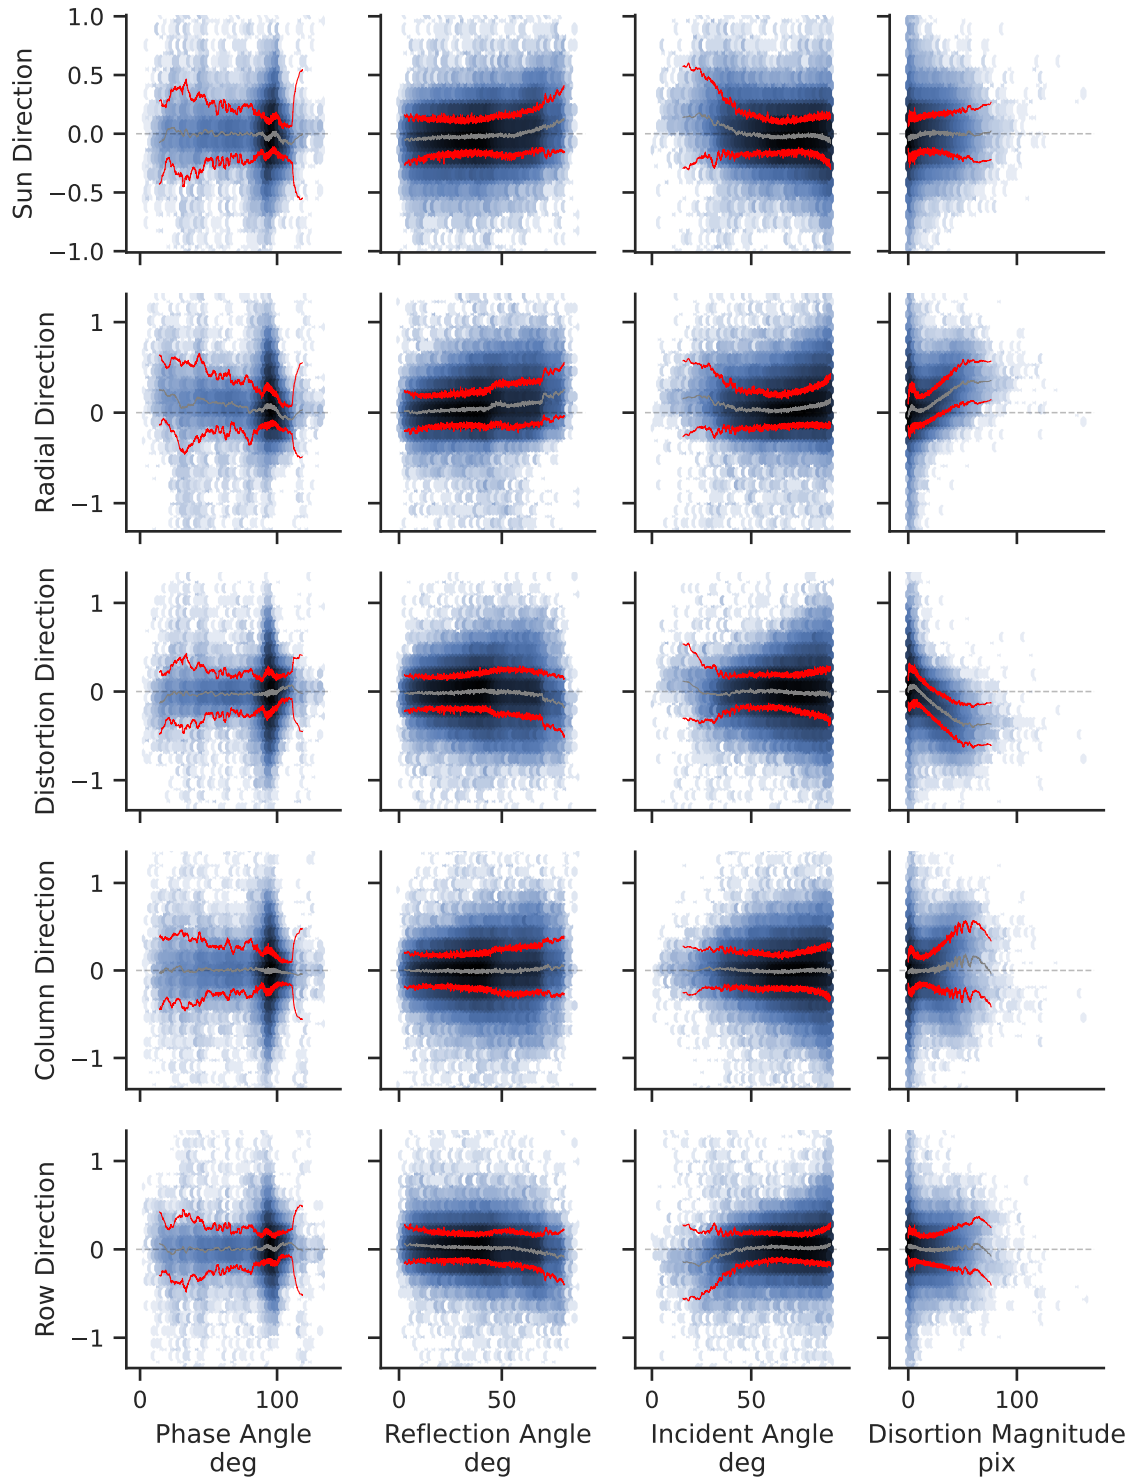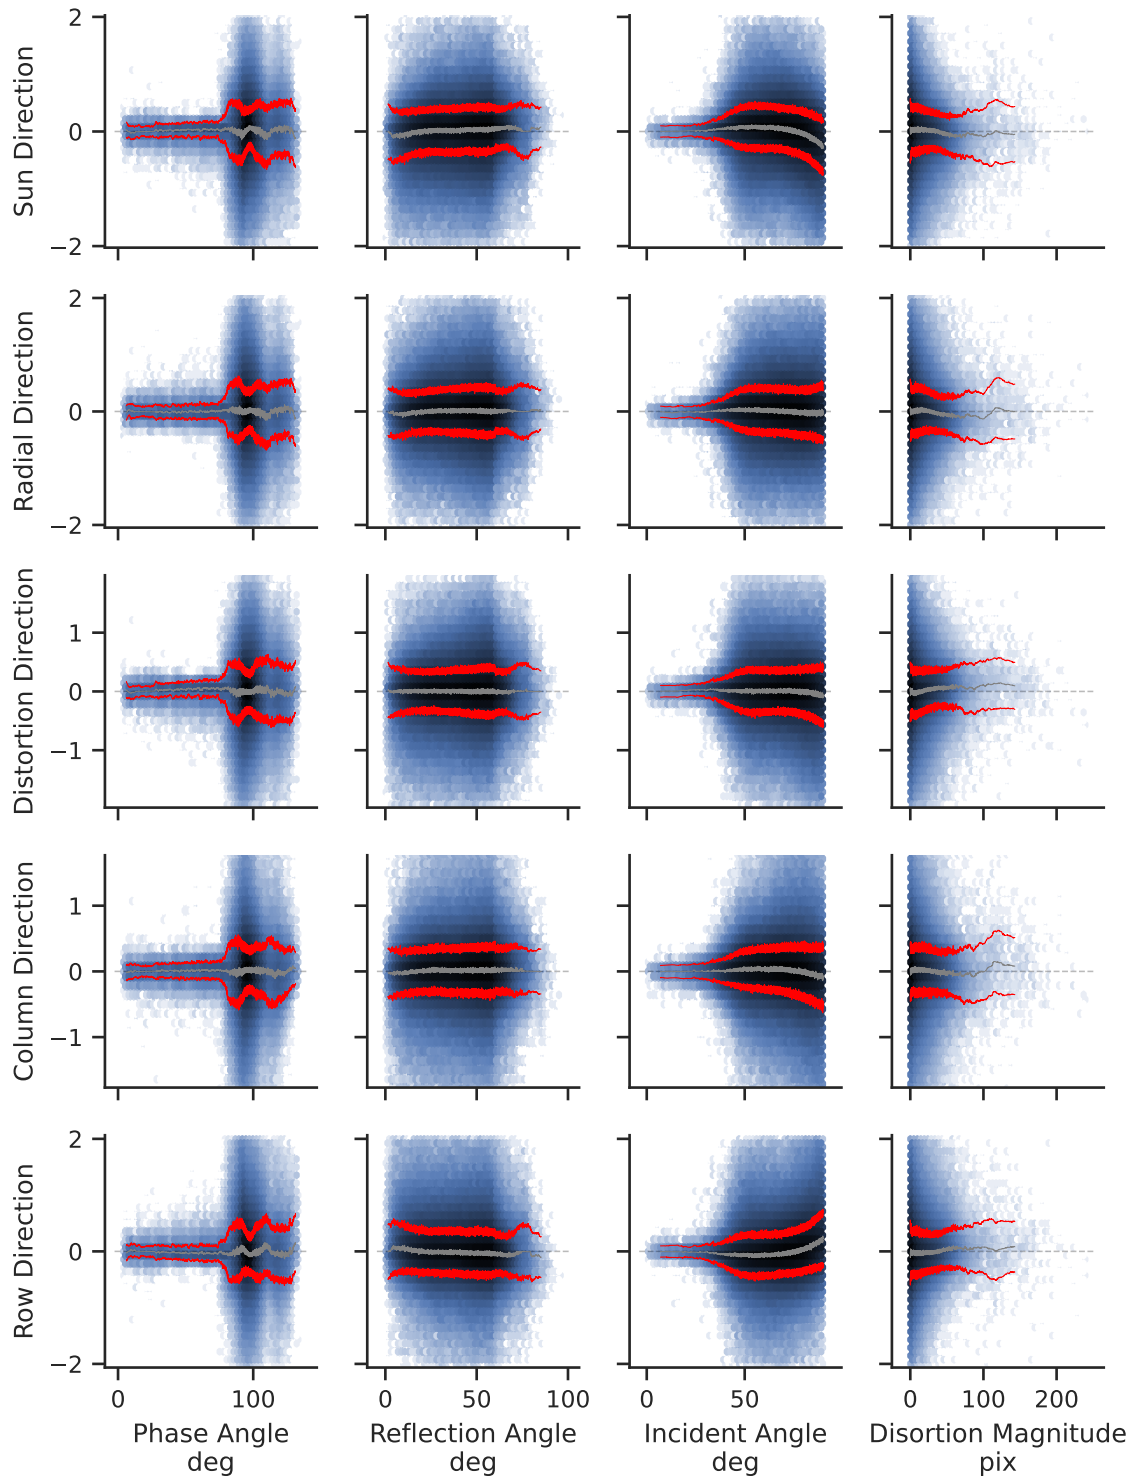
Figure 18. The post-PnP solution observed minus computed residuals in multiple directions against multiple parameters for the GIANT SFN algorithm applied to the OLA V21 model.
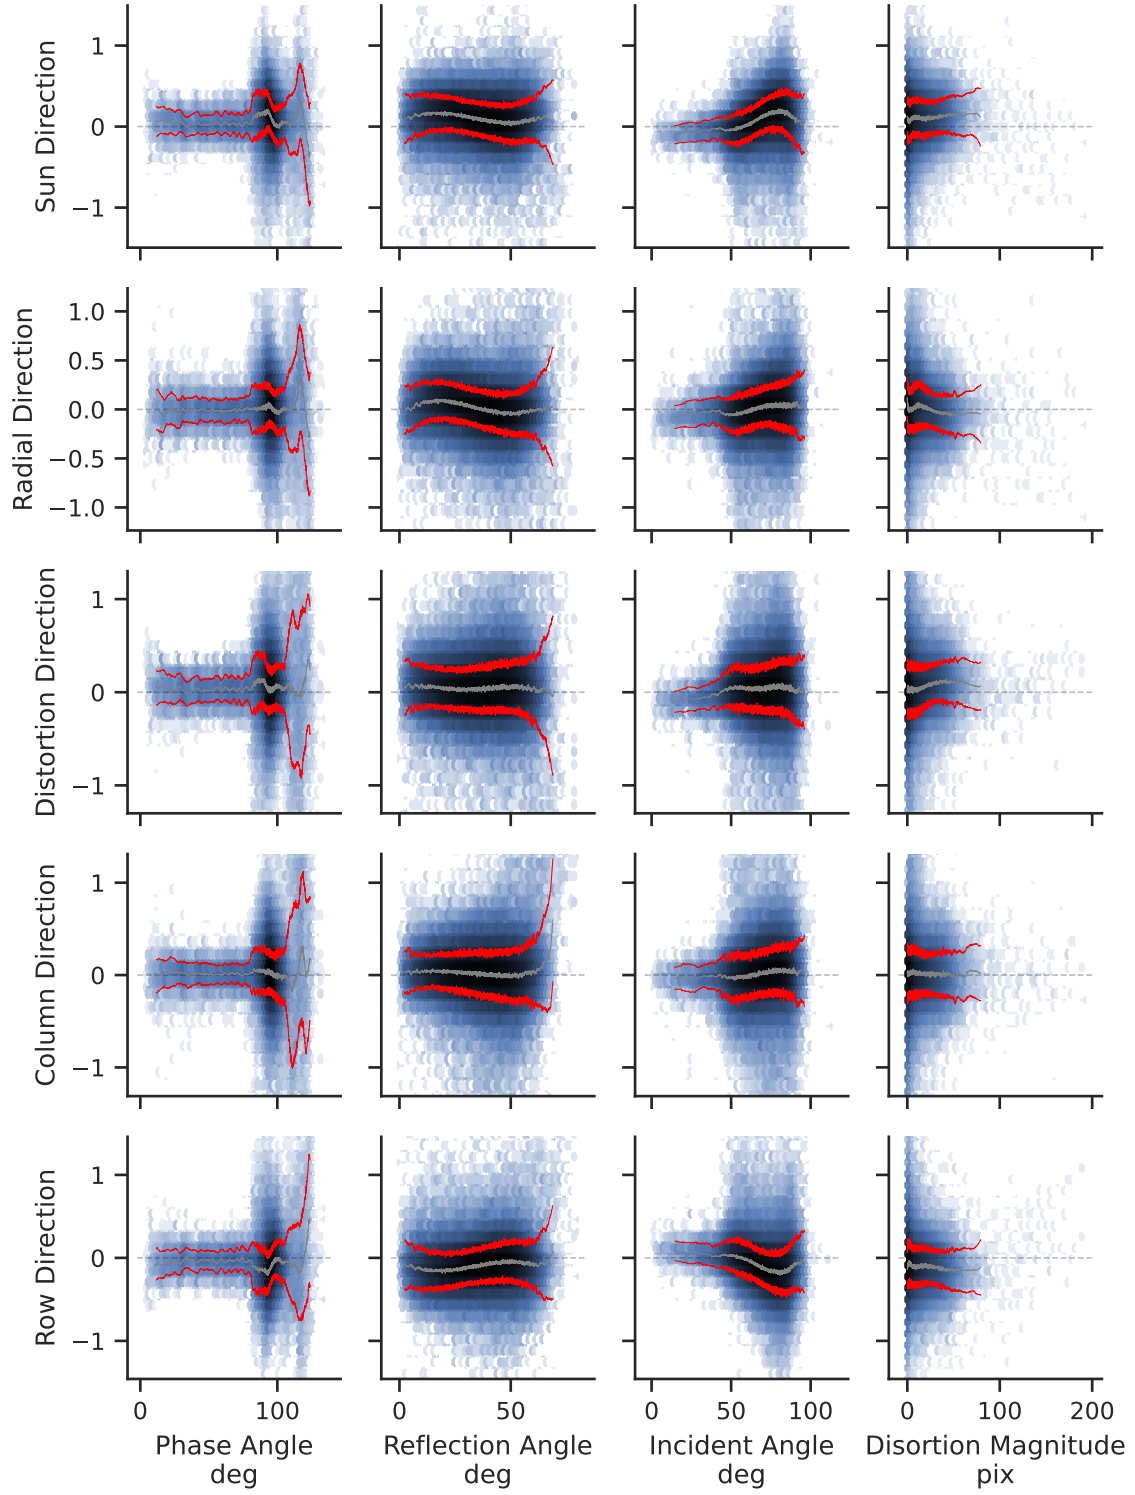
Figure 19.   The post-PnP solution observed minus computed residuals in multiple directions against multiple parameters for the SPC Autoregister algorithm applied to the OLA V21 model.
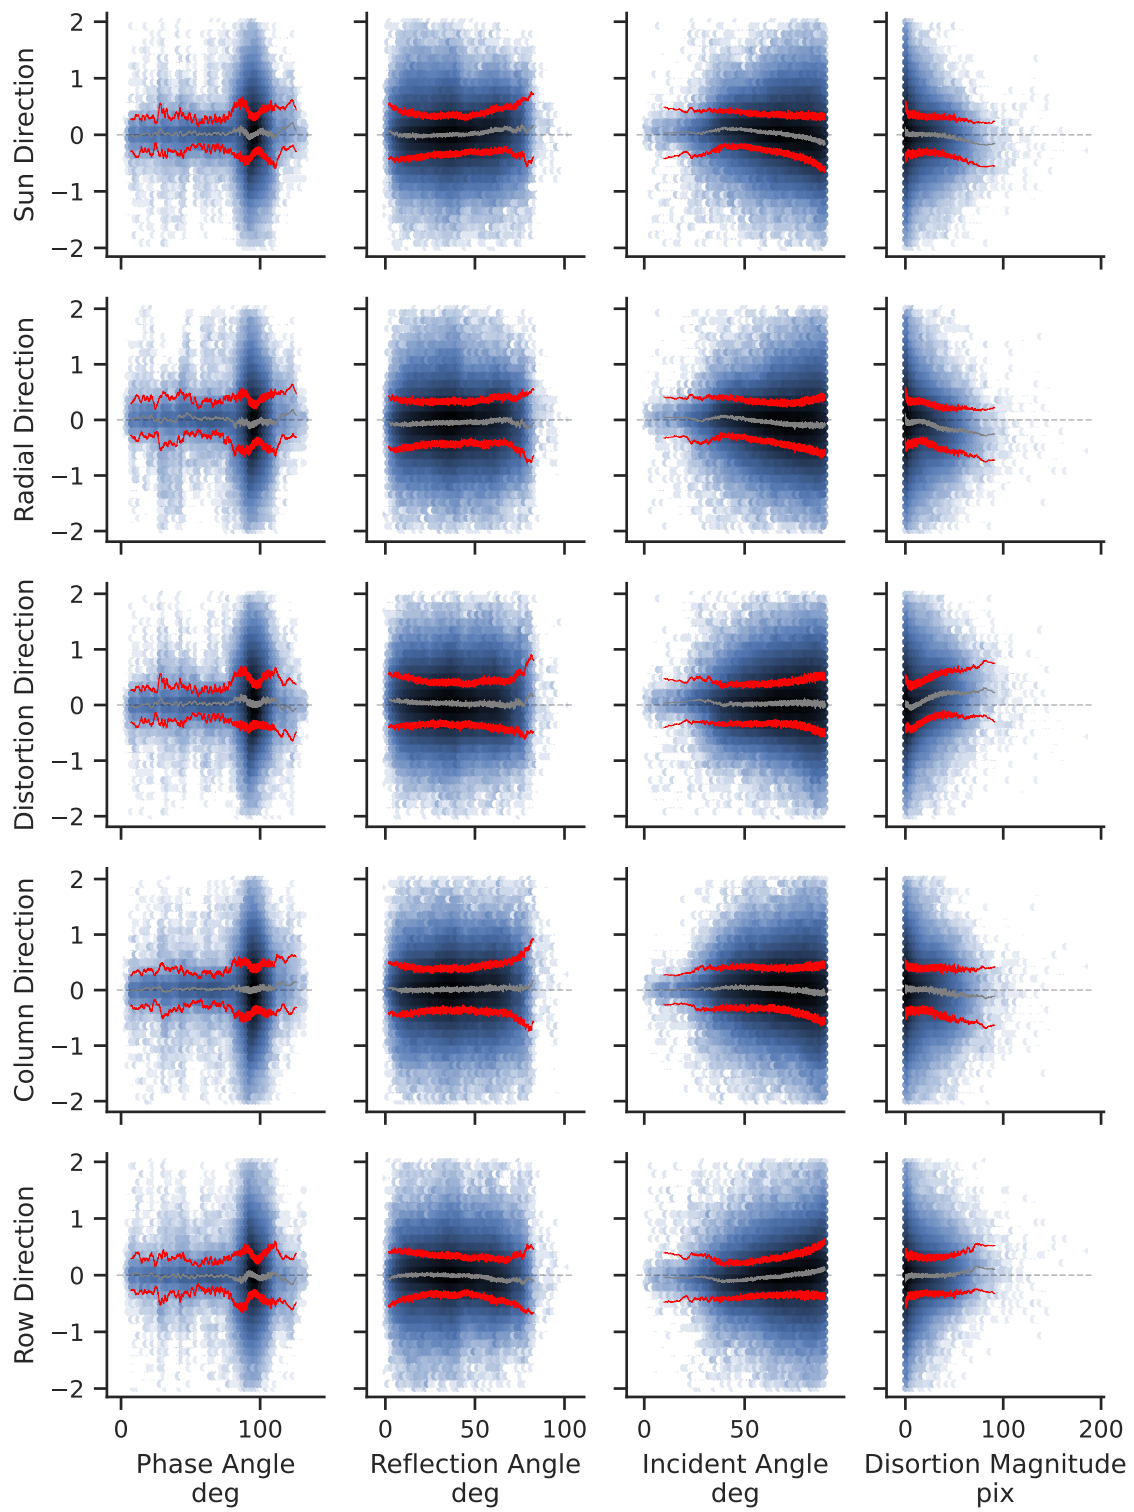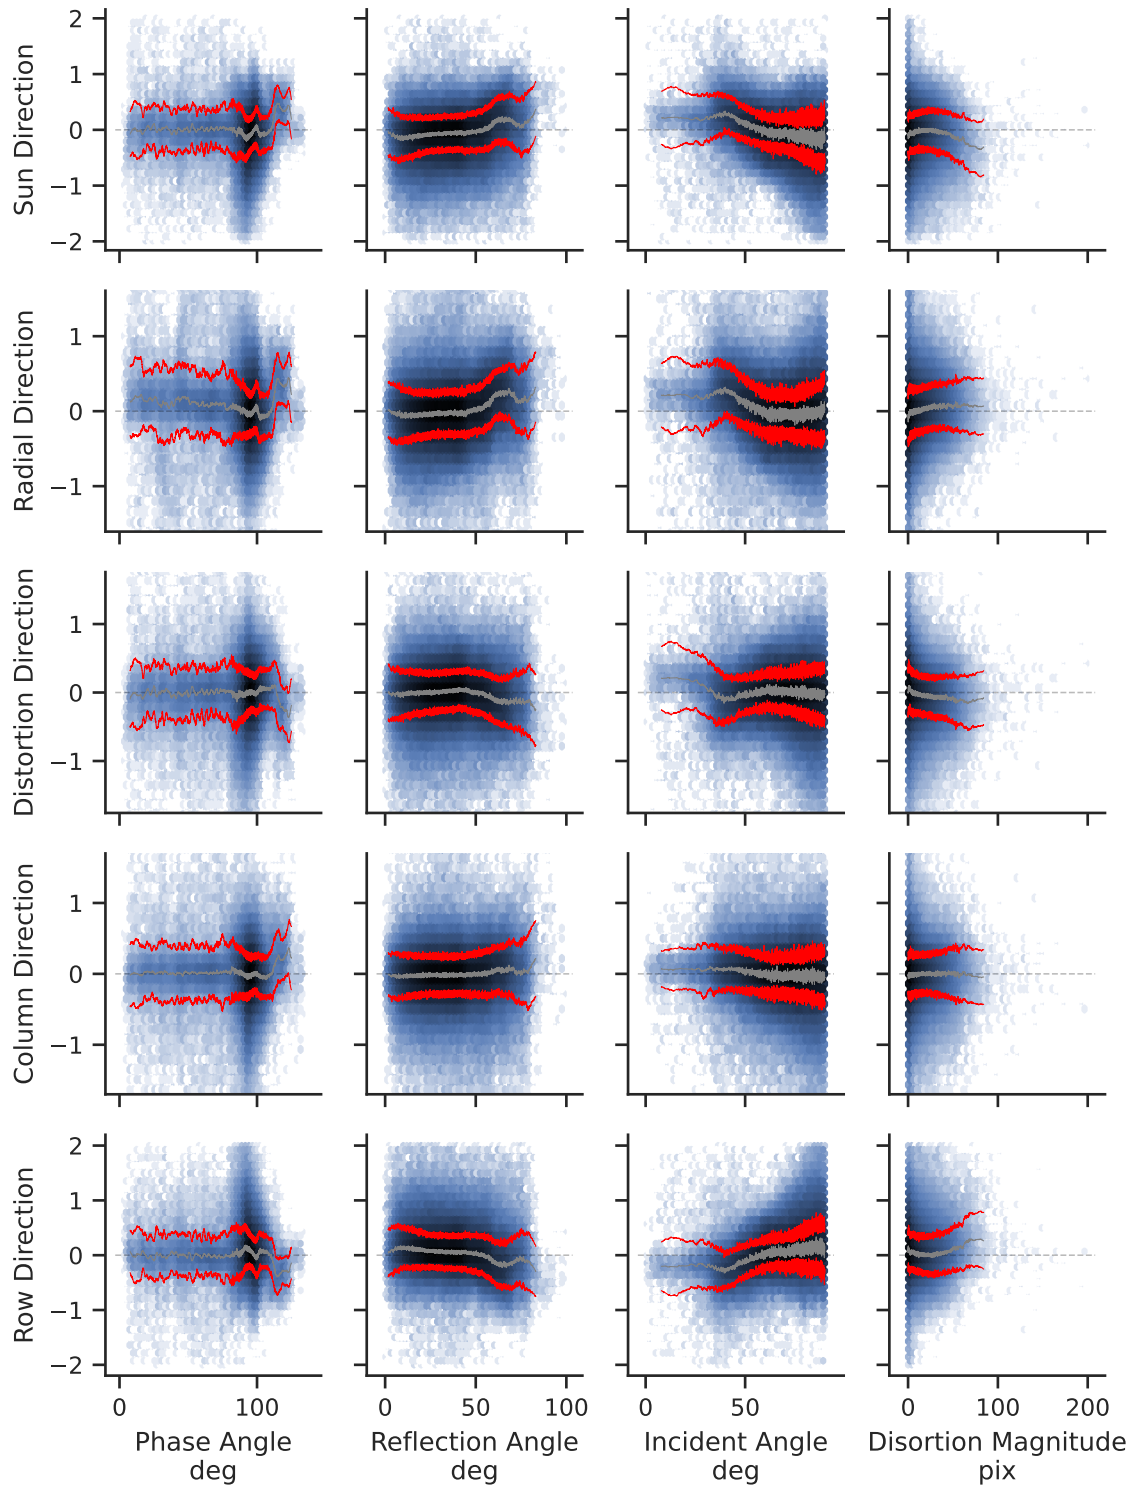
Figure 20. The post-PnP solution observed minus computed residuals in multiple directions against multiple parameters for the GIANT SFN algorithm applied to the SPC V42 model.

Figure 21. The post-PnP solution observed minus computed residuals in multiple directions against multiple parameters for the SPC Autoregister algorithm applied to the SPC V42 model.

Figure 22. The post-PnP solution observed minus computed residuals in multiple directions against multiple parameters for the GIANT SFN algorithm applied to the SPCOLA model.
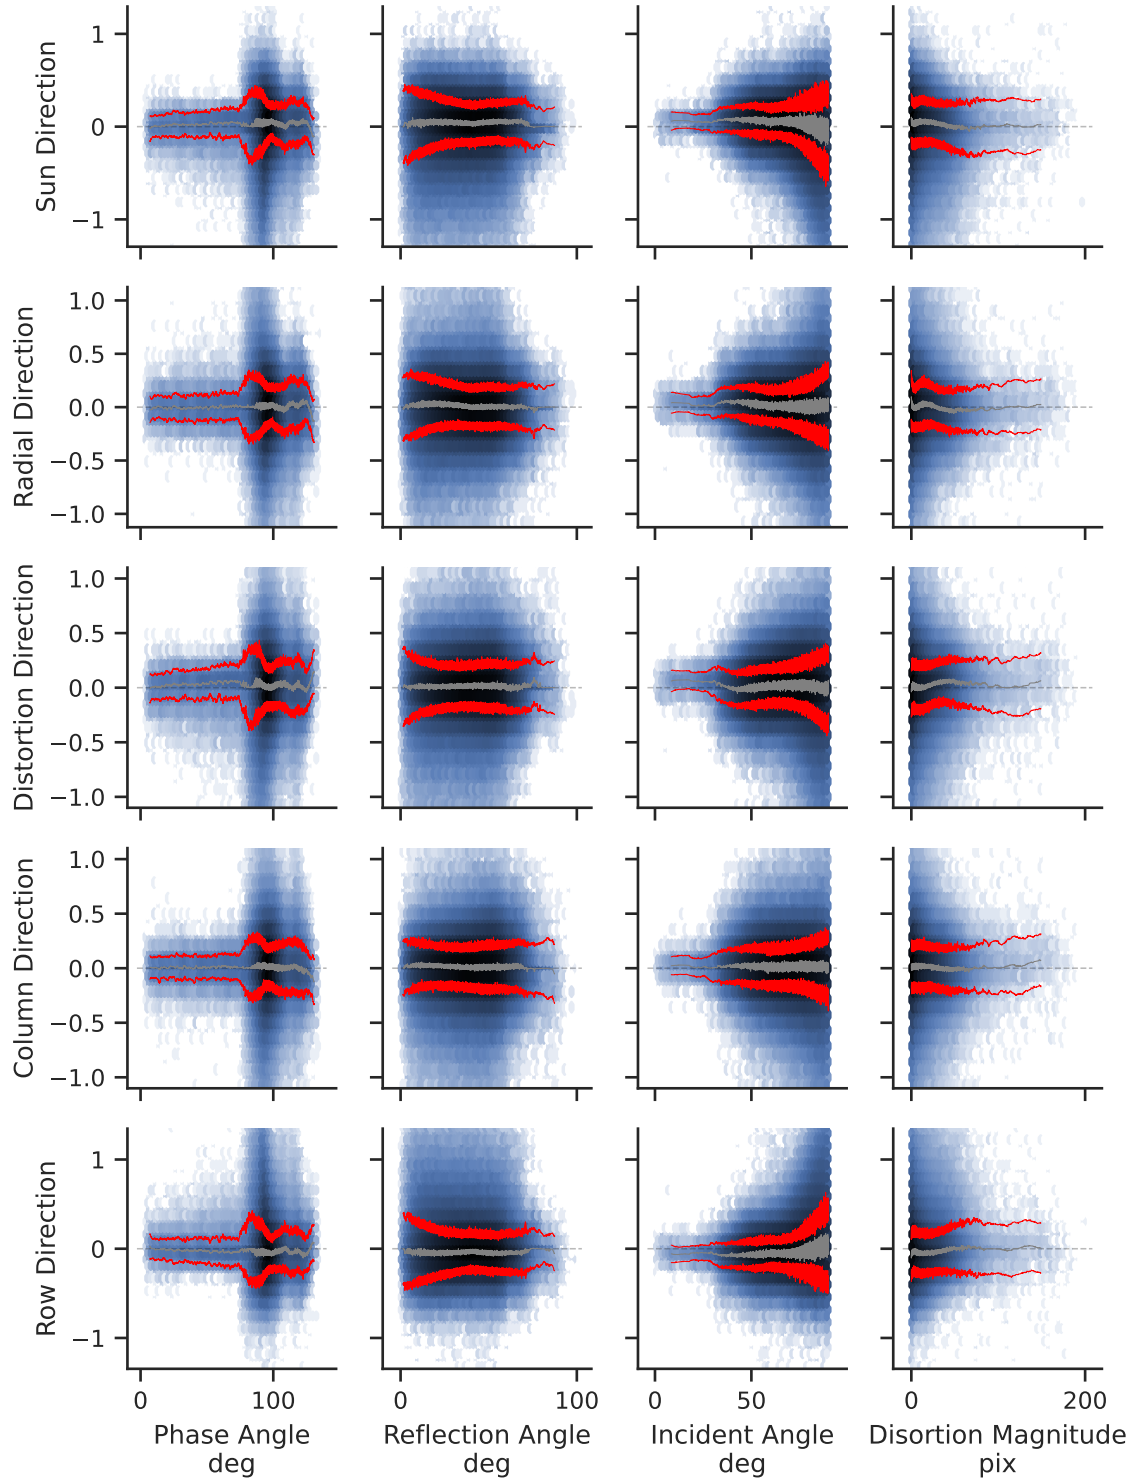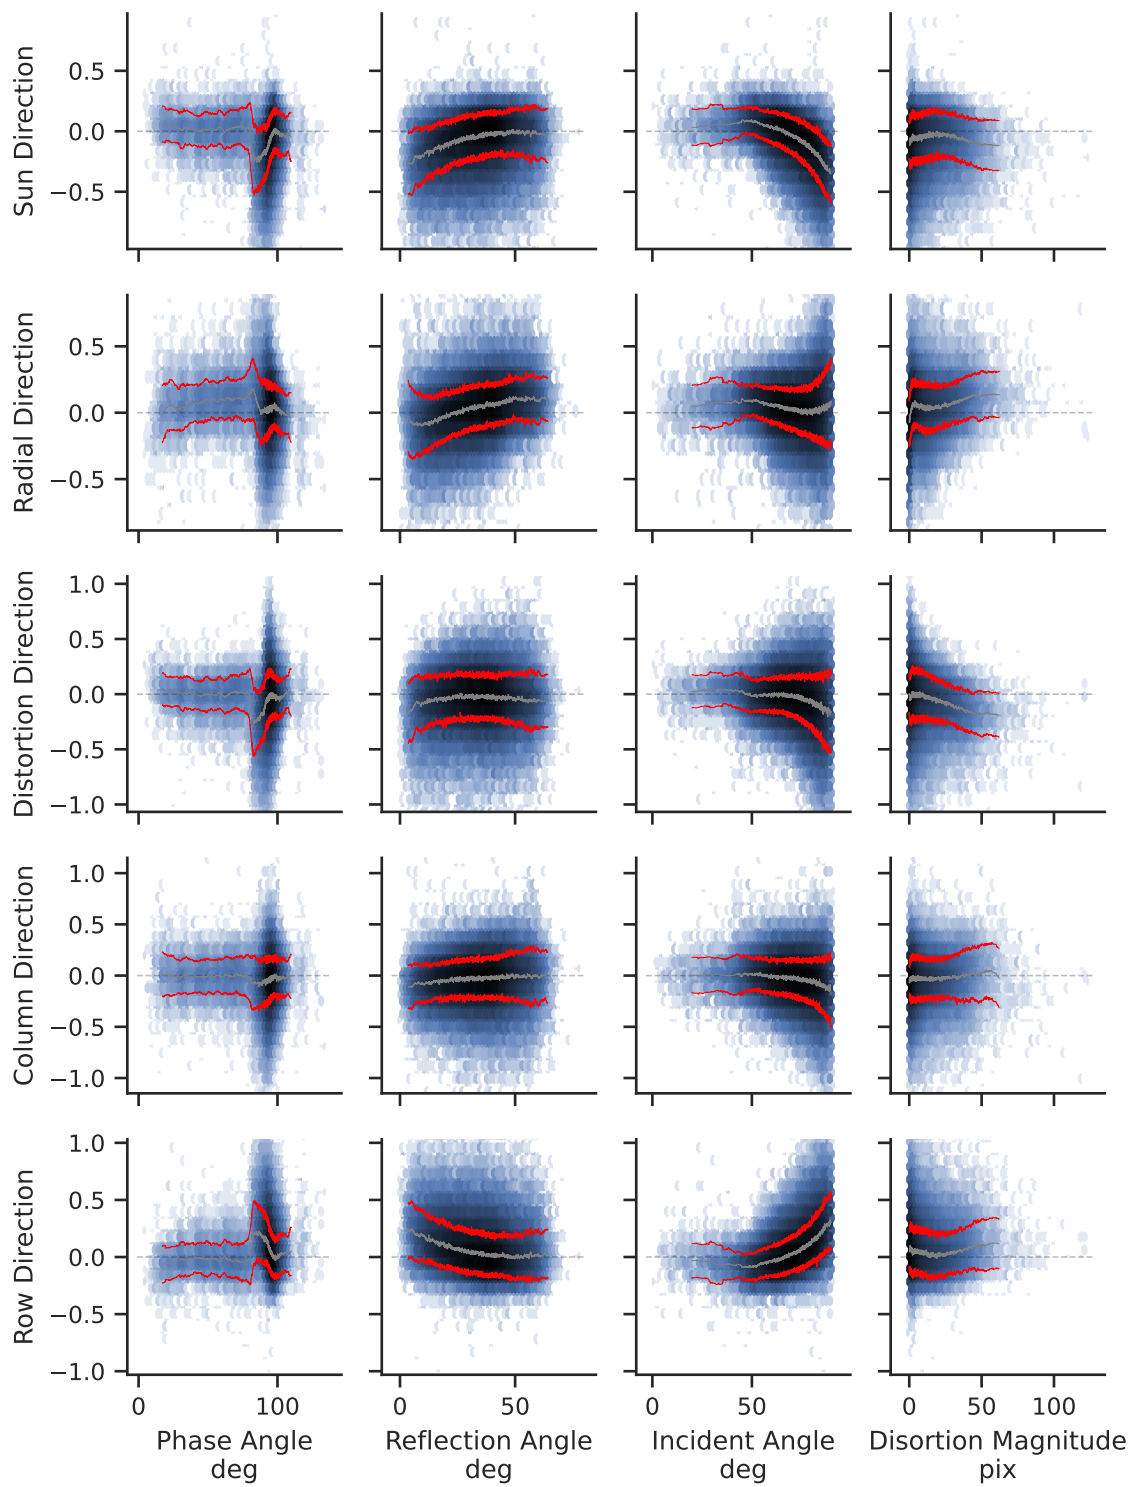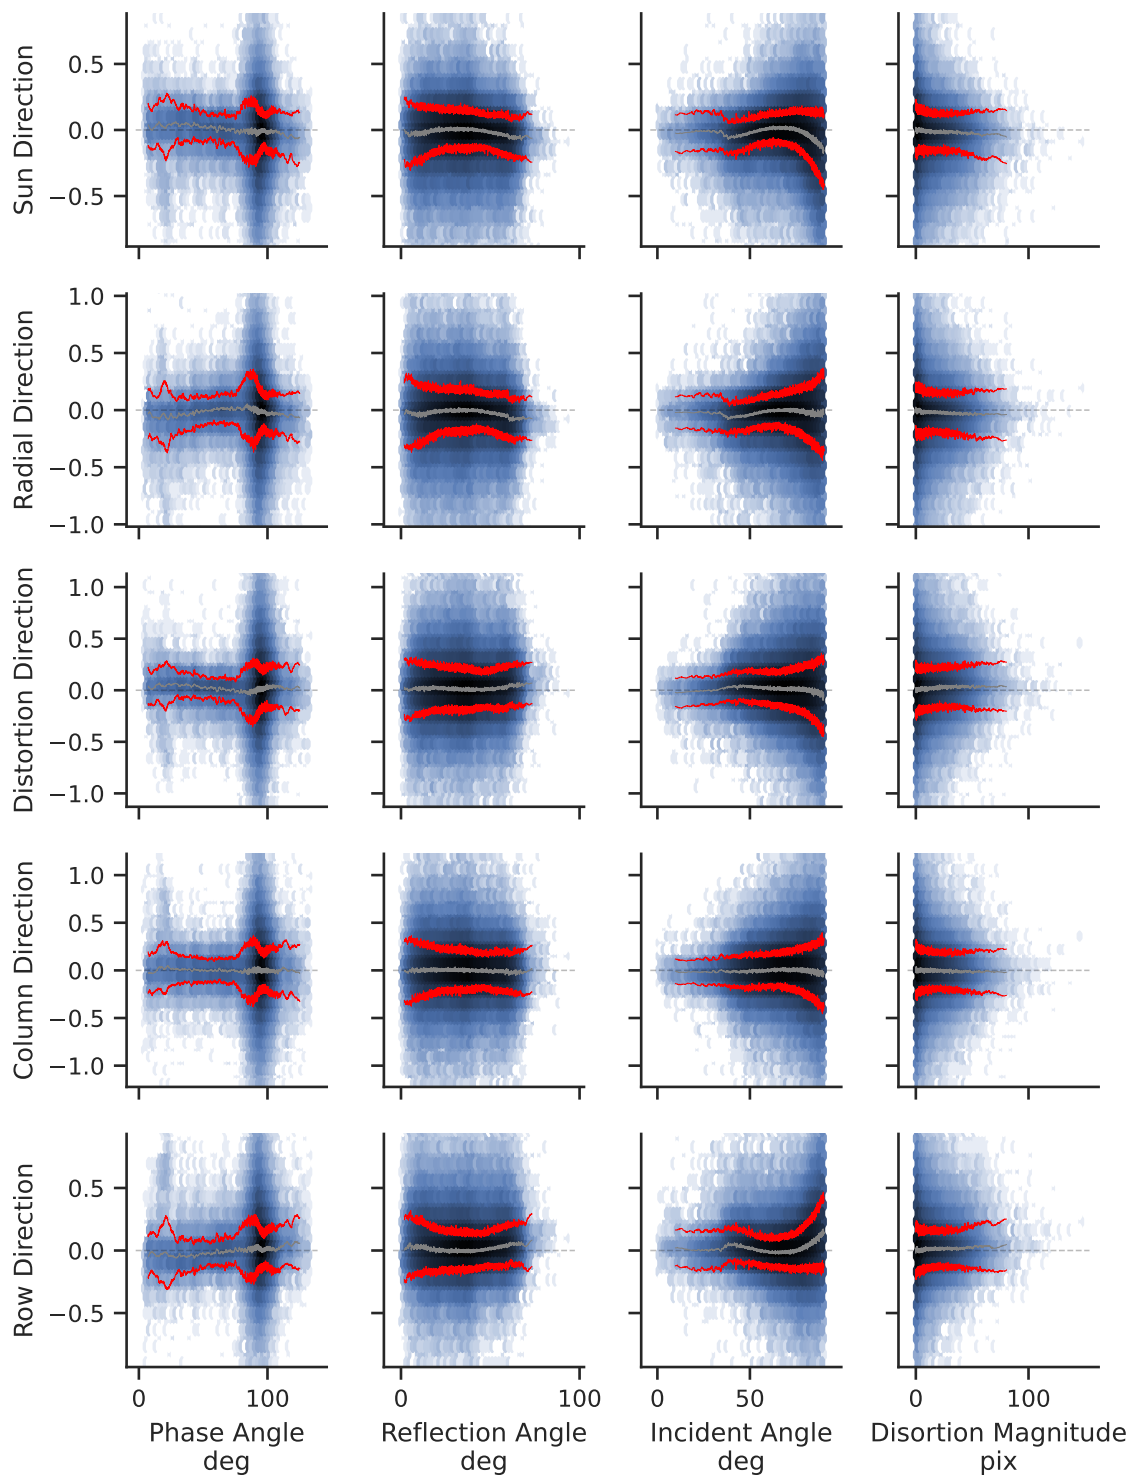
Figure 23. The post-PnP solution observed minus computed residuals in multiple directions against multiple parameters for the SPC Autoregister algorithm applied to the SPCOLA model.

Figure 24. The SPC Autoregister minus GIANT SFN difference for common features in multiple directions against multiple parameters for the OLA V21 model.

Figure 25. The SPC Autoregister minus GIANT SFN difference for common features in multiple directions against multiple parameters for the SPC V42 model.
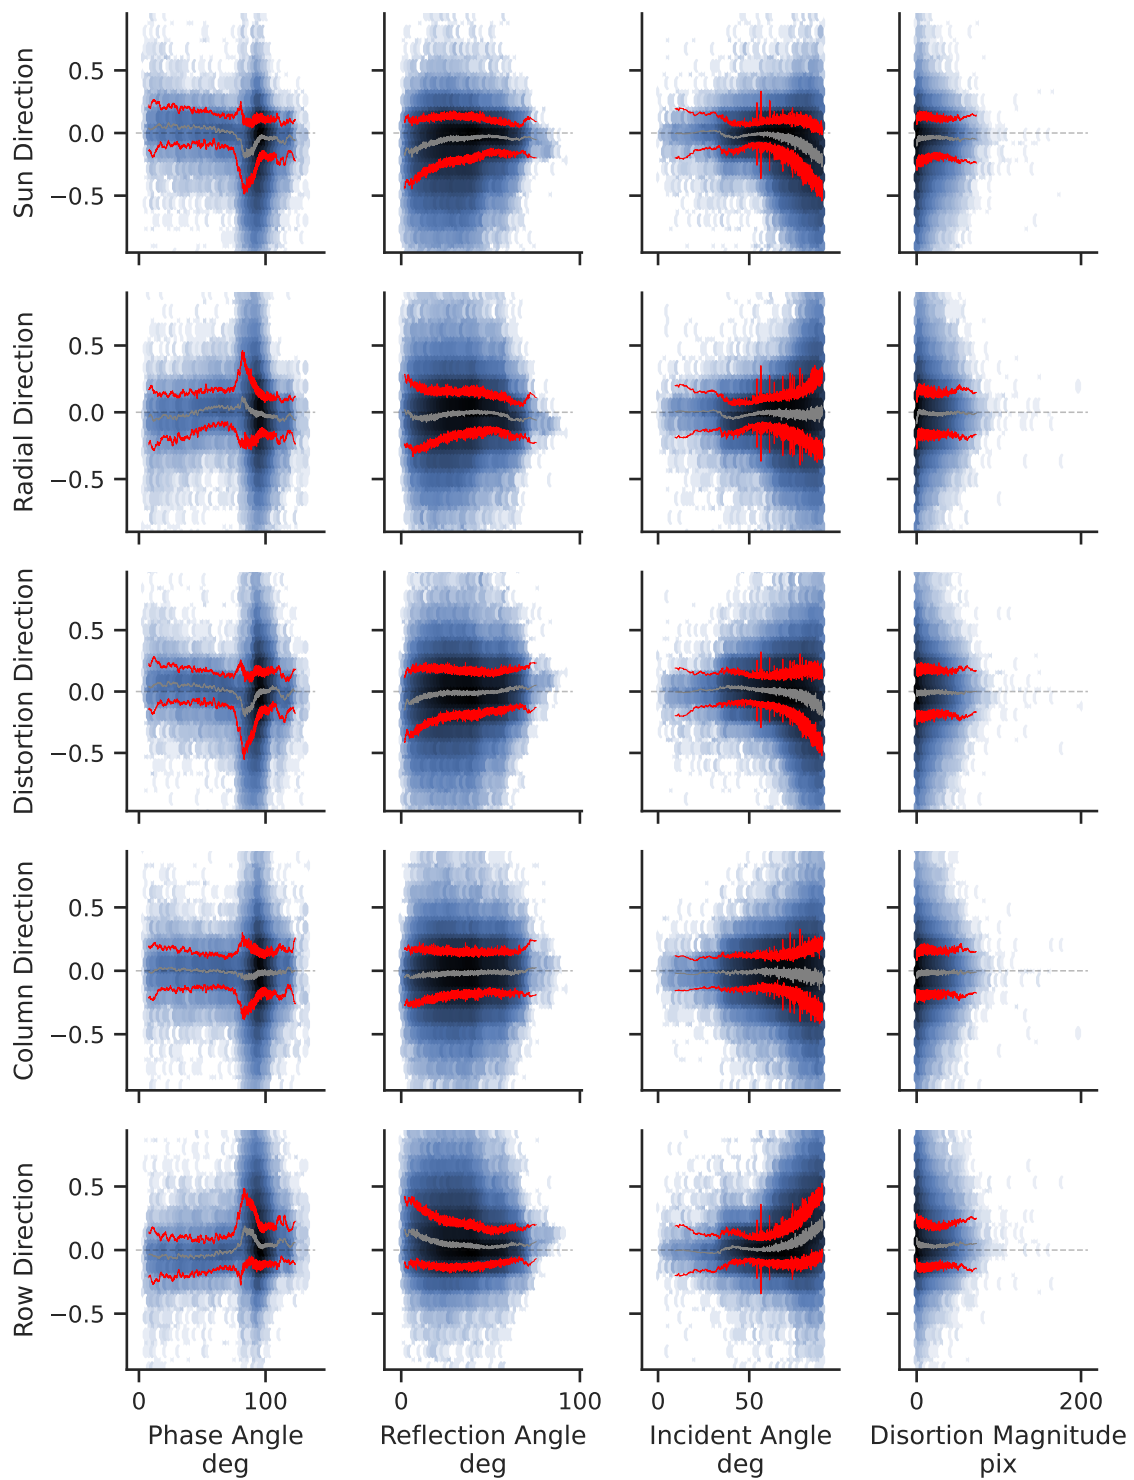
Figure 26. The SPC Autoregister minus GIANT SFN difference for common features in multiple directions against multiple parameters for the SPCOLA model.

TABLES

Table 1. A condensed comparison between SPC Autoregister and GIANT SFN. Modified from Reference 4 with permission.

	SPC Autoregister	GIANT SFN
Prediction	The template is created by using the McEwen BRDF to calculate illumination values for every point on the feature topography.	The template is rendered into the image frame via ray tracing. If the ray does not intersect the maplet, then that pixel is removed from the intersection mask.
Extraction	The image intensity values are projected onto the feature topography to produce an extracted image. A simple method is used to identify shadows and obstructed topography pixels, which are not included in the correlation.	All the pixels in the collected image that are of empty space are included in correlation.
Registration	Normalized cross-correlation between the predicted image and the extracted image takes place on the feature topography. Topography pixels where shadows are expected to be, or pixels in the image detected to be in shadow (due to being below a certain threshold), are ignored in the correlation.	Cross-correlation between the template and image takes place in the image frame. The intersection mask and space mask are combined through a logical OR so that only relevant portions of both the template and image are correlated. Shadows are included in the computation of the correlation surface.
Registration Location	Registration occurs on the feature topography, limiting observable shifts to the size of the feature.	Registration occurs in the image frame. This can accommodate larger initial errors in the *a priori* pose, but can also lead to slight biases due to camera distortions if the *a priori* pose is particularly poor.

Table 2. Tuning parameters for GIANT SFN.

Parameter	Description	Value	units
grid size	Sets the square root of the number of rays per pixel.	1	N/A
minimum correlation score	Sets the minimum cross correlation score for a feature to be considered identified.	0.4	N/A
search region	How far to search around the predicted locations of the feature.	20	pixels
second search region	How far to search around the predicted locations of the feature after correcting *a priori* errors with the PnP solution.	5	pixels
position sigma	The *a priori* uncertainty on the relative position.	0.1	km
attitude sigma	The *a priori* uncertainty on the relative orientation.	0.002	deg
illumination function	The function used to convert the ray trace geometry to intensity values.	McEwen[7]	N/A
run PnP solver	Specifies to run the PnP solver.	True	N/A
PnP RANSAC iterations	The number of RANSAC iterations to attempt when solving the PnP problem to remove outliers.	100	N/A
gsd scaling	The amount that the ground sample distance (gsd) of the camera and the feature can differ for the feature to be possibly visible in the image.	2.0	N/A
reflectance angle maximum	The maximum reflectance angle (between line of sight direction and feature normal vector) for a feature to be possibly visible in the image.	90	deg
incident angle maximum	The maximum incident angle (between sun direction and feature normal vector) for a feature to be possibly visible in the image.	90	deg

Table 3. Tuning parameters for SPC Autoregister.

Parameter	Description	Value	units
minimum correlation score	Sets the minimum cross correlation score for a feature to be considered identified.	0.3	N/A
initial search binning	How much binning to perform when doing the initial search for a landmark (more binning leads to a larger search distance).	5	maplet pixels
final search binning	How much binning to perform when doing the initial search for a landmark (more binning leads to a larger search distance).	1	maplet pixels
position sigma	The *a priori* uncertainty on the relative position.	0.1	km
attitude sigma	The *a priori* uncertainty on the relative orientation.	0.001	deg
illumination function	The function used to convert the ray trace geometry to intensity values.	McEwen[7]	N/A
run geometry solution	Specifies to run the geometry solution (analogous to the PnP solver in GIANT SFN).	True	N/A
invisible limit	The maximum percent of a feature that is considered invisible to the camera for a feature to be possibly visible in the image.	10	percent
maximum emission angle	The maximum emission angle (between line of sight direction and feature normal vector) for a feature to be possibly visible in the image. This is analogous to the GIANT SFN reflectance angle.	90	deg
coverage limit	The minimum percentage of a feature that is illuminated for a feature to be possibly visible in the image.	10	percent
minimum resolution ratio	The minimum ratio between the image ground sample distance and the feature ground sample distance for a feature to be possibly visible in the image.	0	N/A
maximum resolution ratio	The maximum ratio between the image ground sample distance and the feature ground sample distance for a feature to be possibly visible in the image.	3.2	N/A